EVERYTHING IS MADE
OF SOMETHING ELSE
A Practical Guide to Parts, Wholes,
and Why Reality Keeps Leaking
Marcus Qadar

Introduction: Why Parts Matter
What Is a Part?
2. What Is a Whole?
3. Composition Without Mysticism
4. The Ship of Theseus Still Floats
5. The Problem of Arbitrary Objects
6. Scattered Objects and Distributed Wholes
7. Biology: Organisms as Negotiated Wholes
8. Technology: Modular Reality
9. Social Wholes and Moral Buck-Passing
10. Who Is Responsible When Parts Misbehave?
11. States, Institutions, and Abstract Beasts
12. Minimalism vs Maximalism
13. Are Wholes Ever Fundamental?
14. Living in a World of Leaky Boundaries

Everything Is Made of Something Else

A Practical Guide to Parts, Wholes, and Why Reality Keeps Leaking

Introduction: Why Parts Matter

You already believe something strange about reality. You believe that your body is made of cells, but you are not your cells. You believe that your phone is made of components, but breaking it into pieces destroys the phone while leaving the pieces intact. You believe that a government can do things no individual member of that government could do alone. You believe that removing a single grain from a heap of sand leaves you with a heap, but removing them one by one eventually gives you nothing.

These beliefs live quietly in the back of your mind, rarely examined, doing tremendous work. They organize how you think about change, identity, responsibility, and what counts as real. When you argue about whether a fetus is a person, or whether a corporation should

be held accountable, or whether your rebuilt vintage car is still the same car, you are making claims about parts and wholes whether you know it or not.

This is the territory of *mereology*, the study of the relationship between parts and wholes. The name sounds technical because philosophers prefer Greek words for ideas that make everyone uncomfortable. What mereology actually studies is the basic architecture of reality: how things compose, decompose, persist, and disappear. It asks questions that sound simple until you try to answer them. What makes something one thing rather than many things? When does a collection become a unity? Can something be part of itself? If you replace every component of an object, is it still that object?

These questions matter because reality keeps leaking through our categories.

The boundaries we draw around objects turn out to be provisional, context-dependent, sometimes pure fiction. Your immune system contains cells that used to belong to your mother. The forest you walk through is partly fungus living inside tree roots. The computer running this text is executing millions of semi-independent processes that no designer fully understands. We act as if the world comes pre-divided into clean units, and then we're shocked when our tools for thinking about those units fail catastrophically.

The Reductionist Trap

There is a powerful intuition that whispers to us constantly: everything reduces to its smallest parts. This feels true because it's been verified a thousand times over. Chemistry reduces to physics. Organisms reduce to cells. Societies reduce to individuals making choices.

The whole is nothing more than the sum of its parts, and once you understand the parts, you understand everything.

This intuition has given us antibiotics, microprocessors, and the standard model of particle physics. It's not wrong, exactly. It's incomplete in a way that creates enormous blind spots.

Consider a traffic jam. You can describe it completely in terms of individual cars, each following simple rules: accelerate when there's space, brake when there isn't. Every fact about the traffic jam supervenes on facts about individual cars. There is no mystical traffic-jam-substance hovering above the highway. And yet the traffic jam has properties that are nearly impossible to predict from studying any individual car. It emerges from interactions, propagates backward against the direction of travel, and can persist for hours even though no

individual car is stuck for that long. The whole is not more than the sum of its parts, but it's also not just the parts sitting there in a pile.

This is where reductionism becomes a trap. It tells you to look at components when you should be looking at composition. It tells you that wholes are ontologically innocent bookkeeping when sometimes they're the only things doing real work. It whispers that emergence is mysticism when emergence is just what happens when you have parts that interact.

Reductionism fails the moment you try to use it for anything practical. You cannot debug a complex system by examining transistors. You cannot treat a patient by modeling quarks. You cannot reform an institution by changing individual minds. The levels matter. The joints in reality matter. How things hang together matters

as much as what they're made of.

Mereology in Disguise

You have been doing mereology your entire life without calling it that. Every time you fixed something by replacing a part, you made a mereological judgment about which components were functionally independent. Every time you wondered whether your old friend was still the same person after years apart, you were puzzling over identity conditions for persistent wholes. Every time you argued about whether the blame belongs to an individual or a system, you were neck-deep in mereological quicksand.

Engineers do mereology when they design modular systems. Biologists do it when they define organisms. Lawyers do it when they argue about corporate personhood. Political theorists do it when they talk about collective responsibility.

Therapists do it when they help clients distinguish self from family from trauma. The concepts are everywhere, dressed up in domain-specific language, hiding the fact that they're all wrestling with the same deep structure.

This book drags that structure into the light. The thesis is simple, and it will make you uncomfortable: parts and wholes do not behave politely. They shift based on context, purpose, and scale. They leak into each other. They make promises about persistence and break them constantly. Objects are more like knots in a web than atoms in a void, and the web doesn't care about your intuitions.

Once you see this, you can't unsee it. You start noticing the places where reality refuses to be carved at the joints because there are no joints, only different ways of cutting. You start seeing

mereological confusion in arguments that seemed purely empirical or purely moral. You start asking different questions. Not "what is this thing made of?" but "under what conditions does this collection count as a thing at all?"

What This Book Does

This is not a neutral survey. Neutral surveys are for people who want to avoid conclusions. This book has a perspective: mereology is one of the most powerful tools for clear thinking that philosophy has ever produced, and almost nobody knows how to use it outside specialist circles. That's a crime because the questions it addresses are inescapable.

We'll start with the basic machinery: what counts as a part, what counts as a whole, how composition works when it works and how it fails when it doesn't. Then we'll look at the classic problems that have tortured philosophers for

millennia, the ones involving ships and heaps and scattered objects. These problems persist because they're pointing at something genuinely strange about reality, something that can't be solved by being clever with definitions.

Then the book escapes the philosophy department. We'll look at how mereological questions show up in biology, where organisms are constantly negotiating their boundaries with parasites and symbionts. We'll look at technology, where modular design is both the solution and the source of cascading failures. We'll look at social wholes, where responsibility leaks between individuals and institutions in ways that let everyone off the hook simultaneously.

The final sections go deep. We'll confront the question of how many objects actually exist and whether that

question even makes sense. We'll ask whether wholes might sometimes be more fundamental than their parts, a possibility that overturns everything reductionism tells you. And we'll end with the practical upshot: mereology as a tool for noticing where your thinking has gone wrong because you've drawn a boundary in the wrong place and then defended it like your life depended on it.

A Warning

This book will not give you final answers. Mereology doesn't work that way. What it gives you is better questions and a finer sense of where reality is actually carved versus where we've made cuts for convenience. You will leave this book more confused about some things, but confused in a productive way, the kind of confusion that comes from seeing complexity you'd been ignoring.

You will also leave with a suspicion that most confident claims about what things are and how they're organized are papering over mereological chaos. That's uncomfortable. Good. Comfort is what you get from pretending the world has clean edges. This book is about what happens when you look at the edges closely and discover they're more like gradients, negotiations, or outright fictions.

Reality is made of parts. Reality is also made of wholes. These two facts are in constant tension, and that tension is not a bug to be fixed but a feature to be understood. Let's begin.

1.What Is a Part?

Pick up your phone. Now ask yourself: what are its parts?
You might say the screen, the battery, the processor, the case. That's reasonable. But what about the individual transistors in the processor? What about the molecules in the glass? What about the pixels, which aren't physical objects but functional units? What about the software, which makes the hardware behave like a phone instead of an expensive paperweight? What about the cell tower it's currently connected to, without which most of its functionality disappears?

The question "what are its parts?" seemed straightforward. It isn't. Parts are not lying around waiting to be discovered like fossils in rock. They're the result of decisions about what kind of thing you're

analyzing and what kind of boundaries matter for your purposes. This is not postmodern nonsense about reality being subjective. The phone is real, its constituents are real, but which constituents count as parts depends on what you're trying to do.

This chapter is about learning to see that dependence clearly. Because once you understand what makes something a part, you understand why most arguments about composition are arguments about purpose dressed up as arguments about metaphysics.

Physical Parts: The Easy Cases That Aren't

Start with the obvious cases. The leg of a chair is part of the chair. The heart is part of the body. The wheel is part of the bicycle. These feel like objective facts about the world, and in some sense they are. But notice what's doing the work:

spatial enclosure, physical attachment, and functional contribution.

The leg is spatially connected to the chair and holds it up. The heart is spatially inside the body and pumps blood. The wheel is attached to the bicycle frame and enables movement. Take away any of these features and the part-whole relationship starts to dissolve. A detached wheel sitting in your garage is no longer part of your bicycle, even though it was yesterday. The heart, removed during surgery and sitting in a cooler, is no longer part of you in the same way, even though it's still your heart in terms of origin and DNA.

This tells you something important: physical parts are often defined by relationships, not by intrinsic properties. The wheel doesn't have "bicycle-part-ness" written into its atomic structure. It becomes a part by virtue of how it's

connected and what it does. Change the relationship, change the parthood.

Now complicate it. Is the blood inside your heart part of your heart, or part of your circulatory system, or part of your body more generally? Is the air inside your lungs part of you? It's spatially enclosed, it's temporarily under your control, but you'll exhale it in a few seconds. What about the bacteria in your gut, which outnumber your own cells and perform functions essential to your survival? Are they part of you or passengers you're hosting?

The easy cases turn out to have fuzzy boundaries. We draw lines where it's convenient to draw them, and then we forget we drew them.

Functional Parts: Purpose Changes Everything

A different way to think about parts is

functional. Something is a part of a whole if it plays a role in what the whole does. The carburetor is part of the engine because it mixes air and fuel. The editor is part of the newspaper because they decide what gets published. The melody is part of the song because it gives the song its identity.

Functional parts can be physically separate. The moon is not part of the Earth in a compositional sense, it's a distinct astronomical body. But it's part of the Earth-moon system if you're studying tidal forces or orbital dynamics. The relationship is functional, not spatial. The two bodies form a whole for certain purposes and remain separate for others.

This is enormously useful because it lets you track what's doing work in a system without getting distracted by physical boundaries. But it also means that parts can appear and disappear based on

context. The server hosting this document is part of the system that delivered these words to you. Five minutes ago it was part of a different system serving a different request. The physical object didn't change. Its functional role did, and that changed what whole it belonged to.

Functional analysis gets really strange when you apply it to people. Is the night-shift manager part of the factory? Obviously yes in some sense, she makes decisions that keep production running. But she goes home at dawn. Is she still part of the factory when she's asleep in her apartment? What about when she quits and someone else takes the role? The function persists, the person changes, so which one is the part?

You can see the problem. Thinking in terms of functional parts is powerful, but it means accepting that parthood is often

temporary, context-sensitive, and tied to purposes that might be purely human inventions. The factory doesn't care whether it has parts. We care because we're trying to understand or manage it.

Conceptual Parts: Abstractions All the Way Down

Now leave the physical world entirely. What are the parts of a novel? Chapters, obviously. Scenes within chapters. Characters, in some sense, though characters aren't spatially located in the text. Themes, which are even more abstract. The sentence "Call me Ishmael" is part of *Moby-Dick*, but so is the idea of obsession, and so is the genre convention it's responding to.

Conceptual parts are whatever we treat as components when we analyze abstract objects. They're not less real than physical parts, they're just real in a different way. The bishop in chess is part

of the game in that you can't play chess without bishops (or pieces functioning as bishops). But bishops don't exist outside the game. They're conceptual nodes in a rule structure.

Mathematics is built entirely from conceptual parts. A proof has premises, inference steps, and a conclusion. A group has elements, an operation, and axioms it must satisfy. These parts are as well-defined as the parts of a machine, maybe more so, but they're not sitting anywhere in space. They're structural positions in abstract systems.

This matters because much of what we argue about involves conceptual parts. Is consent part of the definition of theft, or is it a separate consideration? Are rights part of personhood or derived from something else? Is the unconscious part of the mind? These questions sound metaphysical but they're really questions

about how to structure concepts in useful ways. The boundaries we draw here determine what counts as a coherent thought.

The Difference Between Containment and Participation

Here's a distinction that will save you from endless confusion: being inside something is not the same as being part of it.

The money in your wallet is inside your wallet. It's not part of your wallet. Take the money out and the wallet is still a wallet, still performs its function, still has all the properties that make it the object it is. The money is contained by the wallet, not composed by it.

Contrast that with the leather that makes up the wallet. Remove the leather and you don't have a wallet anymore. The leather participates in the wallet's

existence in a way the money doesn't. This is the difference between containment and participation, between something being in a whole and being of a whole.

This distinction is doing quiet work all over the place. The passengers are inside the bus, not part of it. The coffee is inside the mug. The thought is inside your head (maybe), but thoughts aren't part of your head in the way neurons are. Failing to track this distinction leads to bizarre conclusions, like thinking that removing passengers disassembles the bus.

But the line isn't always clean. Is the fuel in the tank part of the car? It's consumed to make the car run, it's contained in a purpose-built compartment, but you replace it constantly. What about data in a computer's memory? It's essential to what the computer is currently doing, but it's also completely replaceable and often

transient. At some point containment shades into participation, and you have to make a judgment call.

Pieces vs. Parts: A Crucial Asymmetry

When you break a window, you get pieces of glass. Are those pieces parts of the window? They used to be, in some sense. But calling them parts now feels wrong. They're fragments, remnants, evidence of a whole that no longer exists. Parts are supposed to contribute to a functioning whole. Pieces are what you get when the whole fails.

This asymmetry reveals something deep: parthood often depends on the whole continuing to exist as a unified object. The heart is part of the body as long as the body is alive and organized. Remove the heart, and for a few minutes you still have a body, but soon you have a corpse. The corpse has components, but calling them parts feels like a category error.

Parts serve wholes. Corpses aren't wholes anymore, they're collections waiting to decompose into smaller collections.

You see this in machines too. A laptop that's been disassembled for recycling doesn't have parts anymore. It has materials that used to be parts when they participated in making the laptop a laptop. The transition from parts to pieces happens when the organizing relationship dissolves. And that organizing relationship is what makes a whole a whole in the first place.

When Parts Depend on Intention

Some parts only exist because someone intended them. The handle of a cup is part of the cup because cups are designed to be held. If the same ceramic object emerged from a natural process with no designer, the protrusion would just be a protrusion. It becomes a handle when it

gets enrolled in a purpose.

This makes artforms and artifacts deeply weird from a mereological perspective. The bridge of a song is part of the song because songwriters think in terms of bridges. The climax is part of the story because narrative structure is an intentional construction. These parts are real, they organize the experience, but they're sustained by collective agreement about how to parse aesthetic objects.

Take that agreement away and the parts vanish. Play a symphony for someone from a culture with completely different musical conventions, and they might not hear movements or themes or development. They hear sounds organized by different principles, parsing the experience into different parts. The sounds don't change. The parthood does.

The Hidden Role of Use

Finally, recognize that use determines parts in ways we rarely acknowledge. The keys on a keyboard are parts because they're meant to be pressed individually. But the keyboard as a whole is part of a larger computer system, and that system might be part of a network. The part-whole hierarchy stretches up and down depending on what scale you're operating at.

A carpenter sees different parts in a tree than a biologist. The carpenter sees timber, grain, usable planks. The biologist sees xylem, phloem, bark, root systems. Both are carving the tree into parts correctly for their purposes. Neither carving is the "true" decomposition of the tree into its parts. There is no true decomposition. There are only decompositions relative to explanatory or practical goals.

This doesn't mean anything goes. You

can’t carve reality however you want and expect it to cooperate. The tree doesn’t have a part that’s “all the carbon atoms on Tuesdays.” Some carvings correspond to natural boundaries, functional independence, or causal joints. But even the natural boundaries are multiple and context-dependent. The world is rich enough to support many different true decompositions simultaneously.

The Upshot

What is a part? A component that participates in a whole, where participation means spatial connection, functional contribution, conceptual structure, or some combination of these. Parts are real, but they’re real relative to ways of organizing objects into wholes. Change your purpose, change your scale, change your explanatory goals, and the parts change with them.

This is not a defect in our concepts. It’s a

feature of the world being vastly more complex than any single system of parts and wholes can capture. The trick is knowing which decomposition to use when, and being honest about when you're switching between them.

Most confusion about parts comes from people using different decompositions and thinking they're disagreeing about facts. They're not. They're disagreeing about which facts matter, which boundaries are relevant, which purposes should guide the carving. Those are real disagreements, important ones. But they're not about what's "really" a part. They're about what we're trying to do and what concepts serve those goals.

Once you see this, you stop asking "is X really part of Y?" and start asking "for what purposes should we treat X as part of Y?" That's a better question. It's answerable. And it doesn't pretend that

reality has one correct filing system written into the fabric of space.

2. What Is a Whole?

You know what a whole is. A car is a whole. A person is a whole. The United States is a whole. These examples feel obvious until you start pressing on them, and then they collapse into dozens of difficult questions that nobody agrees on.

Is a whole just a collection of parts? Then a pile of car parts in a warehouse is a car, which it obviously isn't. Is a whole something organized? Then what counts as organization, and who decides? Is a whole defined by boundaries? Then where exactly does the ocean end and the atmosphere begin, given that water molecules are constantly evaporating and condensing across that supposed boundary?

The problem with wholes is that we use the concept constantly while having almost no agreement about what makes

something one thing rather than many things. This matters because nearly every metaphysical, scientific, and political question you can ask depends on what counts as a unified object versus a loose collection.

This chapter is about the different things people mean when they say “whole,” why those meanings are incompatible, and why you need to be able to switch between them depending on what you’re trying to understand.

Wholes as Sums: The Austere View

The simplest account of wholes is brutally deflationary. A whole is just the sum of its parts. Take all the parts, put them together, and you’ve got the whole. Nothing extra. No mysterious unity hovering over the collection. Just parts, all the way down.

This view has the virtue of ontological

cleanliness. It doesn't multiply entities. It doesn't require you to believe in emergent properties or organizing principles or essential forms. The table is just molecules arranged table-wise. You is just cells arranged you-wise. The economy is just people making transactions. Sum up the parts, describe how they're arranged, and you've said everything there is to say.

Philosophers call this mereological nihilism when taken to its extreme: the view that only the most basic parts are real, and everything else is just conceptual bookkeeping. On this view, there are no tables, only particles arranged in table-shaped configurations. There are no forests, only trees (or really, only cells, or really, only atoms). The wholes we talk about are convenient fictions, shortcuts for talking about complicated arrangements of the only

things that actually exist.

This is philosophically tidy and completely useless for living. You cannot navigate the world by thinking only about particle arrangements. You cannot have a conversation about politics without talking about nations and institutions. You cannot debug software by examining transistor states. The wholes might be “nothing over and above” their parts in some metaphysical sense, but pragmatically they’re the only things worth talking about.

Still, the austere view has a point. It reminds you that wholes don’t float free of their parts. Change the parts enough and the whole changes or disappears. Wholes supervene on parts, meaning you can’t change the whole without changing something about the parts or their arrangement. That’s a real constraint, and it prevents you from treating wholes as

magical entities independent of their physical basis.

Wholes as Systems: Organization Matters

A different view says that wholes are systems, and systems are not just sums but structured arrangements where the relationships between parts matter as much as the parts themselves. A pile of bricks is not a house. The same bricks, arranged correctly and mortared together, become a house. What changed? The bricks are identical. The difference is organization.

This view takes structure seriously. It says that some arrangements create wholes with properties that can't be predicted just by listing the parts. A heart pumps blood. A pile of heart cells does nothing. The cells have to be arranged in a specific three-dimensional architecture with the right chemical gradients and

electrical connections before you get pumping behavior. The organization is doing real work.

Systems thinking is everywhere in science because it's how things actually work. Ecosystems have energy flows and feedback loops that can't be understood by studying individual organisms in isolation. Economies have boom-bust cycles that emerge from the interactions between buyers, sellers, and credit systems. Brains have neural firing patterns that constitute thoughts, and those patterns are properties of networks, not individual neurons.

The systems view acknowledges that wholes can have causal powers their parts lack. The traffic jam slows you down. Your heart keeps you alive. The economy crashes. These are whole-level causes, and trying to translate them entirely into part-level language is

possible in principle but useless in practice. You need system-level concepts to track system-level behavior.

But this view has a problem: it threatens to make every collection a whole. Any group of things has some relationships between its parts, even if the relationship is just spatial proximity. If organization is all that matters, then the collection of objects currently on the left side of the planet is a whole, which seems absurd. Systems thinking needs a way to distinguish meaningful organization from arbitrary arrangement, and that's harder than it looks.

Wholes as Illusions: The Deflationary Move

A more skeptical position says that wholes are mostly perceptual or cognitive illusions. We see unified objects because our brains are pattern-recognition machines desperate to

impose order on sensory chaos. The world is actually a buzzing mess of interacting particles, and we chunk that mess into discrete objects because discrete objects are easier to think about and act on.

On this view, the boundaries we draw around objects are psychological, not metaphysical. We say “this is a tree” because treating that cluster of matter as a unified thing helps us predict and navigate the world. But the treeness isn’t out there in nature. Nature just has particles doing physics. We project the category “tree” onto a convenient subset of those particles.

This position gets support from neuroscience. Your brain doesn’t passively record the world. It actively constructs experience from fragmentary sensory data, filling in gaps and imposing structure. The objects you perceive are

the result of sophisticated cognitive processing, not direct access to reality's fundamental furniture. You see wholes because your brain makes you see wholes, not because wholes are independently real.

Taken seriously, this view collapses almost everything you believe. Organisms aren't real, they're just patterns in matter that we find salient. Nations definitely aren't real, they're collective fictions sustained by coordinated behavior. Even your own self might be an illusion, a story your brain tells about a stream of experiences that has no actual unity.

The problem with this deflationary move is that it proves too much. Yes, our categories are partly constructed. Yes, boundaries are often fuzzy or conventional. But some patterns in nature are stable, self-maintaining, and causally

potent independent of what any observer thinks. Atoms were doing their thing long before there were brains to recognize them. Organisms reproduced for billions of years before anyone had the concept “organism.” The patterns aren’t just in our heads.

The Ship, the Swarm, the Crowd, the Corporation

Consider four different kinds of wholes and notice how differently they behave.

The ship is a paradigm case of a material whole. It has clear spatial boundaries, it’s made of physical parts connected in stable ways, it moves as a unit, it has a function (carrying things across water). If you see a ship, you know what you’re looking at. The boundaries are doing almost no work. The ship announces itself.

The swarm is a biological whole without

fixed boundaries. A swarm of bees has a center of mass, moves in coordinated ways, and responds to threats as a unit. But individual bees constantly join and leave. The swarm has no skin, no permanent membership, no clear answer to "which bees are part of the swarm?" It's a whole sustained by behavioral coordination rather than physical attachment.

The crowd is even looser. A crowd at a protest has an approximate location, a shared purpose, and sometimes coordinated action. But the membership is radically unstable. People wander in and out. The edges are completely fuzzy. You can point at the crowd and say "there's a crowd," but you can't say exactly which people are in it versus near it. The whole exists primarily as a social fact, sustained by shared attention and purpose.

The corporation is the strangest of all. It's a legal entity with rights and responsibilities. It can own property, sign contracts, and be sued. It has no physical location in any simple sense. Its "parts" (employees, shareholders, assets) are scattered across the globe, constantly changing, and relate to the corporation through legal and financial abstractions. The corporation persists through complete replacement of every person involved. It's a whole made almost entirely of institutional structure and legal fiction, but it's as real as anything in the economy.

These four wholes show you the spectrum. On one end, you have tight physical integration where the boundaries are obvious. On the other end, you have abstract entities held together by law, agreement, and coordinated belief. Both extremes count

as wholes. That tells you something important: wholeness is not one thing. It's a family of related concepts that shade into each other.

Why "The Whole Is Greater" Is Usually a Lie

People love saying "the whole is greater than the sum of its parts." It sounds wise. It gestures toward emergence, synergy, transcendence. It's also almost always false or confused.

If you mean the whole has properties the parts lack, that's true but trivial. Water is wet. Hydrogen and oxygen atoms are not wet. Wetness is a property of liquid water, not individual molecules. But this doesn't make the whole "greater" than the parts. It makes the whole different from the parts. The wetness emerges from how molecules interact, but it's not surplus to the molecular arrangement. It's just what that arrangement feels like

at the macro scale.

If you mean the whole is somehow more valuable or important than the parts, that's a normative claim masquerading as metaphysics. Sometimes parts matter more than wholes. Your individual life matters more than the corporation you work for, even though the corporation might have more total economic impact. The claim that wholes are greater is often a power move, a way of demanding individual sacrifice for collective goals.

If you mean the whole can't be fully explained by examining parts in isolation, that's true and important. Reductionism fails when you ignore relationships between parts. But this doesn't make the whole "greater." It makes the relationships essential. The whole is not greater than the sum of its parts. The whole is the sum of its parts plus their organization, and organization

is not nothing, but it's also not a mystical bonus.

The slogan persists because it lets people feel like they're saying something profound about emergence without actually explaining anything. Whenever someone invokes it, ask them to specify what they mean. Usually they can't, because they're using "greater" to paper over the hard work of explaining how organization creates new properties.

When a Whole Becomes More Than Bookkeeping

Still, there are cases where wholes do real work that can't be easily reduced to part-level descriptions. The brain is the obvious example. You can describe every neuron's firing pattern, map every synaptic connection, and still not capture what it's like to have a thought. The experiential quality of consciousness seems to be a whole-level phenomenon

that resists decomposition.

Markets are another case. The price of a stock emerges from millions of buy and sell decisions, but once established, that price becomes a causal factor influencing future decisions. The whole (the market price) feeds back on the parts (individual traders) in ways that can't be explained by just summing up individual behavior. You get feedback loops, emergent patterns, and system-level dynamics.

Cultures work this way too. No individual creates a language. Languages emerge from collective usage patterns, but once established, they constrain how individuals think and communicate. The whole shapes the parts that constitute it. This is genuine top-down causation, where system-level properties influence component-level behavior.

These cases matter because they show that wholes can be ontologically

significant without being mystical. You don't need to believe in souls or vital forces to think that some high-level patterns are causally relevant and explanatorily indispensable. You just need to accept that organization creates genuine novelty, and that novelty can matter for how the system behaves.

The Uncomfortable Truth

The uncomfortable truth is that "whole" is not a precise metaphysical category. It's a flexible concept we deploy differently depending on context, and the contexts don't all agree about what makes something genuinely unified versus arbitrarily grouped.

For practical purposes, you treat your car as a whole because it moves as a unit and serves a function. For engineering purposes, you decompose it into subsystems that can be analyzed and optimized separately. For physics

purposes, it's just a temporary configuration of particles. All three perspectives are valid. None is the "true" account of what the car really is.

This pluralism is uncomfortable because we want reality to have a single correct description. We want to know what things really are, independent of our purposes and perspectives. But parts and wholes don't work that way. They're organizational concepts, and organization is always relative to explanatory or practical goals.

The world gives you constraints. You can't organize things however you want and expect them to cooperate. Some carvings correspond to stable patterns, functional independence, causal structure. But within those constraints, there are many valid ways to group things into wholes. The world is complex enough to support multiple true

decompositions simultaneously.

This doesn't make wholes subjective or arbitrary. It makes them purpose-relative. And being purpose-relative is not the same as being unreal. Money is purpose-relative, and it's as real as anything in your life. Laws are purpose-relative. Languages are purpose-relative. These things do work in the world because we collectively treat them as real, and that collective treatment sustains their reality.

The Takeaway

A whole is whatever counts as a unified object for your purposes. Sometimes that's a tightly integrated physical system. Sometimes it's a loosely coordinated collective. Sometimes it's an abstract entity sustained by institutional structure. These are all genuine wholes in different senses, and you need to be able to switch between senses fluidly.

The error is thinking there's one kind of wholeness that everything worth calling a whole must exemplify. There isn't. Wholeness comes in degrees and varieties. Some wholes are robust against perturbation, others fragile. Some have clear boundaries, others dissolve at the edges. Some persist through complete replacement of parts, others are destroyed by minimal change.

What makes something a whole is not meeting some strict metaphysical criterion. It's playing the right kind of role in explanation, action, or experience. Wholes are real because they do work. The ship carries cargo. The swarm forages. The crowd protests. The corporation signs contracts. These are genuine causal processes involving entities that behave as units, even when the unity is partial, temporary, or sustained by human agreement.

Stop looking for the essence of wholeness. Start tracking what different kinds of wholes can and can't do, how they persist, how they interact with their parts, and when treating something as a whole is useful versus misleading. That's the skill mereology teaches. Not "what is a whole really," but "what work does treating this as a whole allow me to do, and what does it obscure?"

That's the question that matters. Everything else is philosophy students arguing about how many angels fit on the head of a pin.

3. Composition Without Mysticism

You have parts. You have wholes. Now the hard question: how do parts compose into wholes without invoking magic?

This sounds like it should be straightforward. Take some bricks, stack them with mortar, get a wall. Take some walls, arrange them correctly, get a house. Composition is just the process of parts coming together to form wholes. Simple.

Except it's not simple at all. The bricks don't become a wall just by being near each other. Pile them haphazardly and you have a pile, not a wall. The walls don't become a house just by existing in proximity. Stand four walls in a field separated by ten meters and you have four walls, not a house. Something has to happen between "parts exist" and "whole exists," and that something is

composition.

The question is what composition actually is. What makes the difference between a mere collection and a genuine whole? What turns many into one? And can you answer this without appealing to mysterious forces, vital essences, or organizational spirits that somehow glue parts together?

This chapter argues you can. Composition is physical, explicable, and everywhere. But it's also more varied and context-dependent than you probably think.

Contact: The Simplest Case

Start with the most basic kind of composition: parts touching each other. Two blocks of wood placed edge to edge compose a longer block (provisionally, for some purposes). The contact is doing the compositional work. Separate them

by a millimeter and you have two blocks, not one.

Contact composition is why you can build things. Lego bricks compose into structures because they interlock. The friction and physical interference between studs and sockets holds them together. Remove the contact and the structure collapses into constituent bricks scattered on the floor.

But contact alone doesn't guarantee composition. Put two strangers next to each other on a subway and they're in contact. They don't compose a larger person. Set two books side by side on a shelf and they're touching. They don't compose a single book. Contact is necessary for physical composition in many cases, but it's not sufficient.

What's missing is the right kind of contact. The Lego bricks are designed to connect. The interface between them is

structured to create stable attachment. The strangers are just bumping into each other. Their contact is incidental, temporary, and doesn't create any functional integration.

This tells you something important: composition requires the right kind of relationship between parts, and "right kind" depends on what type of whole you're trying to create. For artifacts, the right kind usually involves design. For natural objects, it involves physical or chemical bonding. For abstract objects, it involves logical or conceptual relationships. Contact is one way of establishing those relationships in the physical domain, but only when the contact creates integration.

Bonding: When Parts Stick

A stronger form of composition involves actual bonding between parts. Chemical bonds, molecular forces, electromagnetic

attraction. This is how molecules compose into larger structures, how cells compose into tissues, how metals compose into alloys.

Bonding is more robust than mere contact because it resists separation. You can pull two blocks apart easily if they're just touching. You can't pull apart a water molecule without breaking chemical bonds, which requires energy. The bond creates compositional stability.

Steel is an example. It's iron with small amounts of carbon distributed through the crystal lattice. The carbon atoms bond with the iron in ways that prevent the crystal structure from slipping. This makes steel much harder than pure iron. The composition is not just "iron plus carbon." It's iron and carbon bonded in a specific configuration that creates emergent mechanical properties.

Your body uses bonding constantly.

Proteins are chains of amino acids held together by peptide bonds. Break the bonds and you don't have a protein anymore, you have a soup of amino acids. The bonds maintain the structural integrity necessary for the protein to fold into its functional shape. The shape is what does biological work. The bonds make the shape possible.

But bonding has limits as an account of composition. Many wholes are not bonded. A bicycle is not bonded, it's bolted and welded. A school is not bonded, it's a collection of people and buildings held together by institutional structure. Bonding explains chemical composition well. It explains organizational or functional composition not at all.

Dependency: When Parts Need Each Other

A different way parts compose is through

mutual dependency. Each part depends on the others to function, and the interdependency creates the whole. This is how organs compose into organisms, how components compose into machines, how roles compose into institutions.

Your liver depends on your heart to deliver oxygenated blood. Your heart depends on your liver to filter toxins. Neither can function independently (at least not for long). The mutual dependency means you can't have one without the other if you want a living organism. The parts compose because they're locked into a functional web where each part's contribution requires the others.

Machines work this way too. The engine needs the transmission to convert power into motion. The transmission needs the wheels to transfer motion to the ground. The wheels need the engine to provide

power. Remove any component and the machine fails. The dependencies are the compositional glue.

This kind of composition is fundamentally different from bonding. The parts don't need to be physically connected in any strong sense. They need to be functionally integrated. A car's battery can be removed and replaced without destroying the car, as long as some battery is present to play the required role. The composition tolerates part-replacement in ways bonded composition doesn't.

Dependency-based composition also explains why some wholes survive damage better than others. Redundant systems have multiple parts that can perform the same function. Lose one kidney and you still have a functional organism because the other kidney can compensate. Redundancy is insurance

against compositional failure. It means the whole can persist despite losing parts that would otherwise be essential.

Coordination: When Parts Synchronize

Some wholes compose through coordination rather than bonding or dependency. The parts align their behavior, creating patterns that wouldn't exist if each part acted independently. This is how birds compose into flocks, how traders compose into markets, how neurons compose into thoughts.

A murmuration of starlings is the canonical example. Thousands of birds moving in synchronized patterns, creating shapes that flow and morph across the sky. No bird is in charge. No bird knows the global pattern. Each bird follows simple local rules: match the velocity of nearby birds, maintain separation, avoid predators. Those local

rules, executed by thousands of individuals, create the coordinated whole.

The composition here is behavioral, not physical. The birds don't bond or attach. They maintain spacing. The whole exists as a pattern of movement, and the pattern is sustained by each bird responding to its neighbors. Stop the coordination and you have a cloud of individual birds. Maintain it and you have a unified, flowing entity that moves as one.

Markets compose this way too. No trader sees the whole market. Each responds to local information: prices, order flow, news. But millions of individual responses aggregate into market-level phenomena: price discovery, liquidity, volatility clustering, crashes. The market is real. It has properties and dynamics. But it's composed entirely through the coordination of independent agents

following their own logic.

Coordination-based composition is fragile in specific ways. It requires information flow between parts. Disrupt the information and the coordination breaks down. It also requires parts to be responsive to each other. Frozen parts can't coordinate. But when it works, coordination creates wholes that are vastly more flexible than bonded or dependency-based wholes. The parts can come and go freely as long as the coordination pattern persists.

When Composition Fails

Composition is not automatic. Parts can fail to compose even when they're in contact, bonded, dependent, or coordinated. Understanding failure modes clarifies what composition requires.

Contact fails when parts don't align.

Stack blocks at random angles and they don't compose into a stable structure, they topple. The contact is there but the geometry doesn't support integration.

Bonding fails when bonds break. Heat a protein too much and it denatures. The peptide bonds remain, but the shape collapses. The composition that made the protein functional disappears even though the chemical constituents are still present.

Dependency fails when the web breaks. Destroy someone's liver and their heart keeps beating, but the organism is dying. The functional integration falls apart because one critical dependency can't be met. The parts are still there. The whole is dissolving.

Coordination fails when information stops flowing or parts stop responding. A flock scatters when a predator attacks because the local coordination rules

(“flee”) override the global pattern (“stay together”). The compositional failure is instantaneous and total.

These failures reveal what composition actually is: a set of relationships between parts that create and maintain integration. Different relationships create different kinds of wholes with different failure modes. Physical composition fails through separation or bond-breaking. Functional composition fails through dependency-breaking. Behavioral composition fails through coordination loss. Know which kind of composition you’re dealing with and you know how the whole can come apart.

The Hidden Role of Intention and Use

Here’s where composition gets philosophically uncomfortable: for many objects, composition requires intention or use. The parts don’t compose into a whole until someone treats them as a

whole for some purpose.

Take a hammer. The head and handle don't compose into a hammer just by being attached. They compose into a hammer when someone uses the combination to drive nails. The use completes the composition. Before the use, you have a piece of metal attached to a stick. The hammer comes into existence when the assembly gets enrolled in hammering practices.

This is obviously true for functional artifacts. A key is not a key unless it's designed or used to open locks. The same piece of metal, shaped identically, is just a weirdly-shaped piece of metal if it's never used for locks and never intended for that purpose. The composition into "key" requires human purposes entering the picture.

But this extends further than artifacts. Species are arguably composed through

human taxonomic practices. The set of organisms we call “Canis familiaris” is a real biological grouping, but the boundaries are conventional. Wolves and dogs can interbreed. The decision to call them separate species is about human purposes (classification, regulation, understanding) rather than sharp natural boundaries.

Nations are obviously composed through collective intention. The United States is a real entity that does things, but it exists because people collectively treat it as real and act accordingly. Remove the collective intention and you have a geographic region with people living on it. The composition happens in the space of social reality, sustained by coordinated belief and behavior.

This is not to say composition is arbitrary or subjective. The constraints are real. You can’t make a hammer from water

because water can’t maintain the structural integrity required. You can’t make a nation from random strangers who don’t share language or coordination mechanisms. But within the constraints, composition often requires purposes, intentions, or uses that are distinctly human contributions.

Composition Across Scales

One of the most important facts about composition is that it happens at every scale, and the compositional rules change as you move up or down.

Quarks compose into protons through strong nuclear forces. Protons and neutrons compose into atomic nuclei through the same forces. Nuclei and electrons compose into atoms through electromagnetic forces. The compositional mechanisms are completely different at each level.

Atoms compose into molecules through electron-sharing (covalent bonds) or electron-transfer (ionic bonds). Molecules compose into macromolecular structures through weaker forces: hydrogen bonds, van der Waals forces, hydrophobic effects. Again, different mechanisms.

Cells compose into tissues through adhesion molecules and extracellular matrices. Tissues compose into organs through structured arrangements and vascular integration. Organs compose into organisms through coordinated physiological systems. Different mechanisms at every level.

The point is not just that composition is scale-dependent. It's that you cannot use the compositional principles from one scale to explain composition at another scale. Chemical bonding doesn't explain how cells compose into tissues. Tissue

adhesion doesn't explain how people compose into societies. Each level has its own compositional logic.

This is why reductionism struggles. You can in principle reduce higher-level composition to lower-level mechanisms, but doing so loses explanatory power. Describing a corporation in terms of particle physics is possible in principle and useless in practice. The compositional relationships that matter at the corporate level are institutional, legal, and financial. Those relationships supervene on physics, but they're not well-described by physics.

Emergent Organization

Sometimes parts compose into wholes that then organize those very parts in new ways. This creates feedback loops where the whole shapes its own composition.

Your genome is an example. Genes code

for proteins. Proteins regulate gene expression. The proteins produced by some genes turn other genes on or off, creating regulatory networks. The whole (the cell's regulatory network) emerges from the parts (individual genes and proteins) but then controls which parts are active and how they behave. The composition creates organization, and the organization modifies the composition.

Ecosystems do this. Plants produce oxygen and organic matter. Animals consume both. Decomposers break down dead organisms, recycling nutrients. The ecosystem emerges from these interactions, but the ecosystem then shapes which species can survive and how populations balance. The whole constrains its parts.

Cities are the same. People create infrastructure. Infrastructure shapes settlement patterns. Settlement patterns

determine economic opportunities. Economic opportunities attract more people. The city emerges from individual choices but then structures the very choices that compose it. The whole feeds back on the parts continuously.

This kind of compositional feedback is where emergence stops being mysterious. The whole isn't some spooky extra entity. It's the pattern created by part-interactions, and that pattern has causal power because it channels how parts behave. The feedback is real causation, just at a different level of organization.

The Takeaway

Composition is not one thing. It's contact, bonding, dependency, coordination, and feedback, in various combinations depending on what kind of whole you're creating. No single compositional mechanism explains everything.

Physical wholes compose through contact and bonding. Functional wholes compose through dependency. Social wholes compose through coordination and collective intention. Complex wholes use multiple compositional mechanisms simultaneously, which is why they're hard to understand and easy to break.

The mysticism people attribute to composition usually comes from conflating these mechanisms. They see that parts compose into wholes and assume there must be one explanation for how that works. There isn't. Composition is diverse, context-dependent, and scale-relative. Once you accept that diversity, the mysticism evaporates.

What you're left with is the hard work of figuring out which compositional mechanisms are operating in any given case. How are these parts relating to each other? What kind of integration are they

achieving? What would break the composition? Those are answerable questions. They require empirical investigation, not philosophical handwaving.

The world is full of wholes composed from parts. But the composition is always specific, always mechanistic, and always explicable if you look at the right level with the right concepts. Mysticism enters when you expect a single answer to a question that has many context-dependent answers.

Stop asking “how do parts compose into wholes in general?” Start asking “how do these parts compose into this whole, and what maintains or breaks that composition?” That question has answers. They’re different for every case. But they’re answers, not mysteries.

4. The Ship of Theseus Still Floats

There is a ship. Over the years, its planks rot and get replaced, one by one. Eventually every original plank has been swapped out. Is it still the same ship?

This problem is over two thousand years old and philosophers still can't agree on the answer. That should tell you something. Either the question is badly formed, or identity across change is genuinely strange, or both.

Most people's intuition is that yes, it's the same ship. You replace parts all the time without destroying identity. You get a new tire, your car is still your car. You cut your hair, you're still you. Continuous replacement through maintenance seems fine. The ship persists through gradual change.

But now complicate it. Someone collects all the removed planks and rebuilds a

ship from the original materials. Now you have two ships. One is the continuously maintained vessel that's been floating and functioning this whole time. The other is built entirely from the original components. Which one is the actual Ship of Theseus?

Your intuitions probably just crashed into each other. The continuously operating ship has functional and spatiotemporal continuity. The reconstructed ship has material continuity. Both seem legitimate. Neither seems obviously wrong. And yet they can't both be the original ship, because that would mean one ship became two ships through a process that didn't involve splitting.

This is the Ship of Theseus problem, and it's not just philosophical recreation. It's about what makes something the same thing over time when everything about it changes. And that question matters for

everything from personal identity to property rights to warranty claims.

Identity Across Change: The Basic Puzzle

The fundamental issue is that objects change while somehow remaining themselves. You are not the same collection of atoms you were seven years ago. Cells die and get replaced. Proteins denature and rebuild. Your body is in constant molecular flux. And yet you're still you. Same person, different matter.

How is this possible? What makes you continuous across total material replacement? It can't be the matter, because the matter is different. It can't be the exact configuration, because the configuration shifts constantly. It has to be something about pattern, continuity, or causal connection that carries identity forward through change.

This matters legally. If you buy a car and gradually replace every part, do you still own the same car or did you accidentally create a new car? If the new car isn't the same car, when exactly did the transition happen? After the first replacement? After 50% replacement? There's no natural boundary, which means identity is somehow preserved across arbitrary amounts of change.

It matters medically. Your transplanted kidney came from someone else. Is it part of you now? It's functioning in your body, integrated with your biological systems, keeping you alive. But it's got someone else's DNA. Is your body now a composite of two people? Or does the kidney become you by virtue of integration, regardless of origin?

It matters personally. Your childhood self had different beliefs, different memories, different personality traits.

Are you still that person? What makes you the same individual despite psychological transformation? If you suffer amnesia and forget everything, are you still you? If your personality changes completely after a head injury, who exactly survived the injury?

The Ship of Theseus is not abstract philosophy. It's everywhere you look, once you start looking.

Replacement, Repair, and Continuity

The standard solution to the Ship of Theseus is spatiotemporal continuity. The ship that's been continuously floating, continuously functioning, continuously occupying a connected path through space and time, is the real ship. The identity is carried by the continuity of organization and function, not by material sameness.

This works well for gradual replacement.

Replace one plank this year, another next year, and the ship maintains its identity because at every moment it's approximately the same ship it was a moment before. The transitions are smooth. No point marks a clear boundary between "old ship" and "new ship." The ship persists by being continuously almost-identical to its past self.

This is how organisms work. Your cells replace themselves gradually. At any given moment, you're mostly the same organism you were yesterday. The continuity is biological and causal. Your current cells are descendants of yesterday's cells. The pattern of organization persists even as the material substrate turns over.

But continuity has problems. Suppose the ship sinks and sits on the ocean floor for fifty years. Then someone raises it, repairs it with new planks, and sails it

again. Is it the same ship? There's a massive gap in spatiotemporal continuity. The ship wasn't functioning as a ship for decades. Yet intuitively it seems like the same ship, restored.

Or suppose the ship is disassembled carefully, every plank labeled and stored, and then reassembled ten years later in a museum. Same ship? There's perfect material continuity but zero functional or spatiotemporal continuity during storage. The ship existed as a scattered set of labeled planks. Did the ship cease to exist and then come back? Or did it persist through its components?

Continuity matters, but it's not sufficient. Identity seems to tolerate gaps and interruptions under the right circumstances.

Why Your Body Is a Moving Target

Your body presents even stranger

identity problems than ships because biological systems are not cleanly bounded.

Start with the microbiome. You contain trillions of bacterial cells that outnumber your own cells. They're essential for digestion, immune function, and even mood regulation through gut-brain signaling. Are they part of you or inhabitants of you? They have different DNA. They reproduce independently. But you can't survive long without them, and they can't survive outside the environments your body creates.

The usual answer is that they're symbionts, not parts. But the boundary is fuzzy. Mitochondria in your cells used to be independent bacteria that got absorbed billions of years ago. Now they're considered organelles, definitely part of your cells. What's the difference between mitochondria and gut bacteria? Mainly

time and integration. The mitochondria have been integrated so long they can't survive independently. The gut bacteria still can, barely.

So are you a single organism or an ecosystem? The answer depends on what you're trying to explain. For most medical purposes, you and your microbiome are treated as a unit. For genetic purposes, only the cells with your DNA count as "you." For immunological purposes, the distinction matters tremendously because your immune system has to distinguish self from non-self, and it treats gut bacteria as non-self even though they're essential.

Then there's the problem of cells that aren't yours. Pregnant women carry cells from their fetuses that can persist for decades after birth. You might have cells from your mother circulating in your body right now. Are those cells part of

you? They're in your body, integrated with your systems. But they're not your genetic lineage.

The boundaries of your body are not where your skin is. They're complicated, negotiated, and context-dependent. You are a moving target because the parts that compose you are constantly being replaced, supplemented, and lost. Identity persists through this flux, but it's not because you maintain material sameness. It's because you maintain organizational continuity and causal connection.

Personal Identity as a Mereological Problem in Disguise

The deepest version of the Ship of Theseus is about persons. You persist through time despite radical change. What makes you the same person you were ten years ago?

One answer: psychological continuity.

You remember being that person. Your current mental states are causally connected to past mental states through a chain of memory and experience. The person is the stream of consciousness, and identity is carried by the continuity of that stream.

This answer immediately hits problems. You don't remember most of your past. Infant you had almost no psychological similarity to adult you. If you develop dementia and forget everything, do you cease to exist? What if you suffer a personality-altering brain injury? The psychological continuity might be completely severed, yet intuitively you still exist, just changed.

Another answer: bodily continuity. You're the same person because you're the same organism, and organisms persist through cell replacement as long as the biological processes continue. Your body

carries your identity forward regardless of psychological changes.

But bodies change too. Lose a limb and you're still you. Receive a transplant and you're still you. In principle you could replace every organ except the brain and still be you. So is the identity in the brain? Maybe, but brains change too. Neurons die, connections rewire, structures atrophy or grow. The brain is not a static object preserving identity through material sameness.

The best answer is probably that personal identity is carried by a combination of causal continuity, psychological connectedness, and bodily continuity, with different factors mattering in different contexts. For legal purposes, bodily continuity often dominates. For personal and moral purposes, psychological continuity matters more. For medical purposes, biological

continuity is key.

But notice what this means: personal identity is not one thing. It's a cluster of overlapping criteria that usually track together but can come apart. When they come apart, we have to decide which criteria matter for the question we're asking. Identity is not metaphysically fundamental. It's a useful concept that works in normal cases and gets messy in edge cases.

The Problem Gets Worse: Fission Cases

The Ship of Theseus involves parts getting replaced. Personal identity gets worse when persons might split.

Imagine your brain is damaged on the left hemisphere. Doctors remove it and transplant the right hemisphere into a new body. The resulting person has half your brain, many of your memories,

much of your personality. Are they you?

Now imagine doctors save both hemispheres and transplant them into two different bodies. Each resulting person has half your brain and claims to be you. Both have causal continuity with you. Both have psychological connectedness. Both have some bodily continuity (the brain matter). Which one is you?

Standard answers collapse here. If continuity matters, both are you. But you can't be two people. If neither is you, then you ceased to exist despite your brain hemispheres surviving and generating conscious experiences. That seems wrong. If only one is you, which one? There's no principled way to choose.

The usual philosophical move is to say that identity doesn't survive fission. You cease to exist and two new people come into being, each with a strong claim to

being your continuation. Identity is not preserved, but something important (your mental life, your memories, your values) is preserved twice over.

This answer is disturbing because it means identity is not about intrinsic facts but about extrinsic facts like whether splitting occurs. If only one hemisphere is transplanted, you survive. If both are transplanted, you die. But from your perspective before the surgery, nothing about you determines which outcome happens. Your survival depends on what happens to parts of you after separation.

Fission cases reveal that identity might not be the deep metaphysical fact we think it is. What's deep is continuity, causation, and psychological connectedness. Identity is a useful label we apply when those factors converge on a single individual. When they diverge, identity breaks down, but the underlying

factors don’t.

Artifacts vs Organisms vs Persons

Different kinds of objects have different identity conditions, and that matters for how they tolerate change.

Artifacts are defined by function and use. A ship is a ship as long as it functions as a ship. Replace every part and it’s still a ship because shipness is about function, not material. Destroy the functional organization (burn it to ash) and the ship ceases to exist even if the atoms are still there.

Organisms are defined by biological continuity. A cat is a cat as long as the biological processes that constitute cat-life continue. Replace cells and it’s still a cat. Stop the biological processes (death) and you have a corpse, not a cat. The material might be identical for a few minutes after death, but the organism is

gone.

Persons are defined by psychological continuity and/or bodily continuity, depending on context. You're you as long as there's an appropriate causal connection between current and past mental states, or as long as your body maintains organizational continuity. Lose memories and you're still you (usually). Lose bodily continuity through fission and identity gets weird.

These different identity conditions mean different objects tolerate different kinds of change. You can repair an artifact extensively without threatening its identity. You can't repair an organism past a certain point because death is identity-destroying for organisms. You can manipulate a person psychologically quite a bit and they remain the same person, but radical changes (complete amnesia, personality inversion) raise

genuine questions about identity.

The Ship of Theseus seems paradoxical because we expect one answer to work for all objects. There isn't one answer. There are multiple answers, each appropriate for different kinds of things. Ships persist through functional continuity. Organisms persist through biological continuity. Persons persist through psychological and/or bodily continuity. Apply the wrong identity condition and you get paradoxes.

When Identity Becomes Conventional

Sometimes identity is just conventional, a matter of what we decide rather than metaphysical fact.

Is the restored ship the same ship if it sat on the ocean floor for a century? There's no fact of the matter. We can stipulate either way. For historical purposes, it might matter that it's the original

material. For maritime purposes, what matters is whether it's seaworthy. For insurance purposes, it's probably a different ship because it was totaled and then rebuilt. Same object, different identity depending on context.

Is the corporation the same corporation after every employee is replaced? Legally, yes. Corporations have legal identity that persists regardless of personnel changes. Culturally, maybe not. A company with completely different people might have a completely different culture and mission. Is it really the same company or just inheriting the name and legal structure?

Is the rock band the same band after the lead singer dies and gets replaced? Fans will fight about this. Some say the essential member is gone, so it's a different band using the old name. Others say the band is defined by its music and

continuing mission, so it's the same band. There's no fact that settles this. It's about what you think carries band-identity.

These cases show that identity is often underdetermined by the facts. The physical and historical facts are what they are, but those facts don't uniquely determine identity. We impose identity on objects based on what matters to us: legal continuity, functional continuity, material continuity, personal attachment, or conventional designation.

This doesn't make identity fake or arbitrary. The conventions are doing real work. They let us track objects through time, assign responsibility, enforce contracts, and maintain relationships. But it means identity is often a practical category rather than a metaphysical one. The world gives us persistence, causal continuity, and organizational structure.

We use those to construct identity concepts that serve our purposes.

The Upshot

The Ship of Theseus still floats because we want it to. We care about tracking ships through time, so we construct identity conditions that allow ships to persist through repair and replacement. The same ship continues because continuity of function, location, and causal history is enough to preserve shipness.

But the ship also reveals that identity is not simple. Objects can change radically while remaining themselves. They can split or merge in ways that make identity ambiguous. They can cease to exist and then arguably come back. Identity is carried by continuity, but continuity comes in many forms: material, functional, causal, psychological, legal. These forms don't always align.

When they don't align, we have to decide what matters. And that decision is often contextual, purpose-relative, and conventional. There's no single fact about identity waiting to be discovered. There are many overlapping facts about continuity and change, and we use those to construct identity concepts that work for the problems we're trying to solve.

The philosophical lesson is to stop asking "is it really the same thing?" and start asking "what kind of continuity matters for this context?" That question has answers. They're different for ships, organisms, persons, corporations, and rock bands. But they're answers, and they let you navigate change without pretending objects have magical essences that persist regardless of what happens to them.

Identity is real. It's also negotiated, contextual, and sometimes conventional.

That combination is what lets you replace every cell in your body and still be you, fix your ship and still own the same ship, and wonder whether the band is really the band without the original guitarist. The wonder comes from expecting one kind of answer when the world gives you many.

5. The Problem of Arbitrary Objects

Right now, there is an object composed of your left shoe and the Eiffel Tower. It exists. It has mass, spatial extent, and a perfectly definite location (scattered across France and wherever you are). You've never thought about it before, but according to certain accounts of composition, it's as real as your shoe or the Eiffel Tower individually.

This sounds insane. Why would such an object exist? What holds it together? What makes it one thing rather than two things that happen to be mentioned in the same sentence? And if this object exists, then so does every other arbitrary combination of things: your coffee mug plus Saturn's rings, the letter Q plus all left-handed people, the concept of justice plus this comma.

Welcome to the problem of arbitrary

objects, where metaphysics collides with common sense and both come away damaged.

The Unrestricted Composition Principle

The problem starts with a seductive idea called unrestricted composition. It says: for any collection of objects, no matter how scattered or disparate, there exists a whole composed of exactly those objects. Always. No exceptions.

Pick any set of things. Those things compose a whole. The whole might be useless, unnamed, and never thought about, but it exists. Composition happens automatically whenever you have multiple objects. You don’t need contact, bonding, organization, or purpose. You just need objects, and composition follows.

This principle has a major advantage: it

avoids vagueness. If composition is restricted, you need criteria for when composition occurs. Those criteria will have borderline cases. Is this pile coherent enough to count as a heap? Is this group coordinated enough to count as a team? Vague criteria mean vague answers, and vague answers are philosophically unsatisfying.

Unrestricted composition eliminates vagueness by saying composition always occurs. No borderline cases. No judgment calls. Any plurality of objects automatically composes a whole. The ontology is maximalist, but it's clean.

The principle also has theoretical elegance. Mereology, the formal logic of parts and wholes, works smoothly with unrestricted composition. You can define operations like "the sum of these objects" without worrying about whether the sum exists. It always exists. You can build a

complete theory without exceptions or special cases.

But the cost is obvious: you have to accept that your left shoe plus the Eiffel Tower is a real object. Along with infinite other combinations that seem absurd.

Do "the Left Half of My Desk Plus Napoleon's Hat" Exist?

Let's make this concrete. Napoleon's hat (assuming it still exists in some museum) and the left half of your desk compose an object. Call it "Deskapoleon." Deskapoleon has mass equal to the sum of its parts. It has a scattered spatial location spanning your home and a museum in France. It has properties: partly wooden, partly fabric, partly historical artifact, partly furniture.

Does Deskapoleon exist? You've never referred to it before. It serves no purpose.

It has no causal unity. The parts don't interact. They're not coordinated. They're not even on the same continent. What would it mean for this to be a real object rather than just a linguistic construction?

The unrestricted composition theorist says: Deskapoleon exists just as surely as your desk exists. Your desk is also just a collection of parts. The fact that desk-parts are spatially proximate and functionally integrated is pragmatically important for you, but metaphysically irrelevant. Existence isn't sensitive to human purposes or pragmatic salience. Either composition produces objects or it doesn't. If it does for your desk, it does for Deskapoleon.

This answer is consistent, but it has implications. If Deskapoleon exists, so does every other arbitrary fusion. There are uncountably many objects occupying

the space you think of as “your desk.” There’s the desk itself. There’s the desk-minus-one-atom. There’s the desk-minus-a-different-atom. There’s the left half. The right half. The top three-quarters. The arrangement of atoms on prime-numbered coordinates. On and on, infinitely.

Most of these objects are useless to think about. We don’t name them, track them, or care about them. But if unrestricted composition is true, they’re as real as the objects we do care about. The world becomes ontologically bloated with objects no one has ever conceptualized or will ever use.

Why Philosophy Keeps Inventing Monsters

Philosophers didn’t invent arbitrary objects to be perverse. They emerged from trying to solve a genuine problem: where do you draw the line?

If you deny unrestricted composition, you need criteria for when composition occurs. When do parts compose a whole versus remaining a mere plurality? The criteria can't be "whenever we find it useful to recognize a whole," because that makes composition mind-dependent. Before humans existed, did composition not occur? That's implausible.

So you need objective criteria. Maybe composition requires contact. But contact is vague. Molecules touch through electromagnetic forces that extend infinitely, just very weakly at distance. Where's the boundary between touching and not-touching? Any line is arbitrary.

Maybe composition requires bonding. But bonding is also vague. Chemical bonds have different strengths. Van der Waals forces are extremely weak bonds. Are they enough? If not, where's the cutoff? Again, arbitrary.

Maybe composition requires functional integration. But function is observer-relative. A pile of bricks has no function. Arrange them into a wall and they function to divide space. Did the bricks compose a whole only when we arranged them? That makes composition depend on human purposes.

Every proposed restriction on composition either admits vagueness, involves arbitrary cutoffs, or makes composition observer-relative. None of these options is philosophically satisfying. Unrestricted composition avoids all three problems by saying composition always occurs. The cost is ontological promiscuity. The benefit is theoretical cleanliness.

This is why philosophers keep inventing monsters. The monsters (arbitrary fusions, scattered objects, weird composites) are the price of avoiding

vagueness and arbitrariness. Whether that price is worth paying depends on how much you care about ontological parsimony versus theoretical simplicity.

Common-Sense Mereology vs Unrestricted Composition

Common sense says composition is restricted. Not every collection of objects composes a whole. Your shoe and the Eiffel Tower don't compose anything. They're just two separate objects that someone mentioned together.

Common sense tracks objects that matter: things with causal unity, functional integration, or practical salience. Your desk is an object because it's causally unified (parts move together), functionally integrated (designed to support work), and practically salient (you interact with it daily). Deskapoleon has none of these features, so it's not an object.

This matches how we actually think and talk. We don't refer to arbitrary fusions. We don't act as if they exist. Our best scientific theories don't posit them. Chemistry talks about molecules, not molecules-plus-random-distant-atoms. Biology talks about organisms, not organism-plus-unrelated-objects. Physics talks about fields and particles, not arbitrary scattered collections.

But common sense has a problem: it can't explain where it draws the line without invoking vagueness or mind-dependence. When exactly does a collection of parts become causally unified enough to compose a whole? How much functional integration is required? Different people will draw the line differently, which suggests there's no objective fact about where it belongs.

Unrestricted composition avoids this by accepting everything. Common sense

avoids it by ignoring edge cases and pretending the boundaries are clearer than they are. Both strategies work until you press on them. Then common sense admits it's using fuzzy concepts, and unrestricted composition admits it's committed to absurd objects.

The Cost of Saying "Yes to Everything"

Accept unrestricted composition and you get some strange consequences beyond just ontological bloat.

First, nothing ever comes into existence or goes out of existence. When you build a chair, you're not creating a new object. The atoms that compose the chair already existed, and according to unrestricted composition, they already composed the chair-object. What you did was arrange them in a way that makes the chair-object salient and useful. But it existed before you built it, just in a scattered, non-

functional configuration.

This seems backwards. Chairs come into existence when you build them and cease to exist when you burn them. But if unrestricted composition is true, the chair-object is eternal. It exists whenever its constituent atoms exist, regardless of how they're arranged. You didn't create the chair, you just made a pre-existing object functional.

Second, objects constantly overlap. The left half of your desk is an object. The right half is an object. The whole desk is an object. The desk-minus-one-atom is an object. All these objects occupy overlapping regions of space. They share most of their parts. They're not separate in any normal sense, they're different ways of carving up the same underlying matter.

This is metaphysically weird. We think of objects as discrete, non-overlapping

entities. But unrestricted composition gives you massive overlap. Every region of space contains infinite objects, all composed of slightly different selections of the available parts. The world becomes a crowded space of overlapping fusions.

Third, you lose the distinction between objects and mere pluralities. Normally we think "some things" and "a thing composed of those things" are different. Three apples in a bowl are three objects (a plurality), not one object (a fusion). But unrestricted composition says there's also a fourth object: the fusion of the three apples. The plurality is many. The fusion is one. But both are equally real.

This distinction is hard to maintain when composition is unrestricted. If any plurality automatically composes a fusion, then talking about pluralities without fusions becomes impossible.

Every time you refer to multiple objects, you've also referred to the fusion of those objects. The ontology collapses the many and the one together in confusing ways.

Restricted Composition: The Alternative

The alternative is to restrict composition: composition only occurs under certain conditions. When conditions are met, parts compose a whole. When they're not, you have a mere plurality.

The obvious question is what conditions matter. Different theories give different answers.

Contact theories say composition requires parts to touch. This rules out scattered objects like Deskapoleon but allows ordinary physical objects. The problem is defining "touch" precisely. Molecules never truly touch, they repel

electromagnetically. So contact is already an idealization.

Bonding theories say composition requires parts to be bonded. This works for molecules but fails for larger objects. Your organs aren't bonded to each other, yet they compose you. Bonding is too restrictive.

Organization theories say composition requires parts to be organized. This captures functional wholes but makes composition observer-relative. Organization is defined relative to purposes, and purposes are mind-dependent. Did organisms not compose before minds existed to recognize organization? That's implausible.

Life theories say composition occurs for living things but not artifacts or aggregates. This gives you organisms as genuine wholes without arbitrary fusions. But it's hard to justify the restriction.

Why would composition care about life? What's metaphysically special about biological organization versus other kinds?

Each restricted view captures some intuitions but leaves others unsatisfied. None avoids all the problems that motivated unrestricted composition in the first place. They trade ontological bloat for vagueness, arbitrariness, or mind-dependence.

The Pragmatic Middle Ground

Maybe the right answer is pragmatic. Composition is unrestricted metaphysically, but salience is restricted pragmatically. All fusions exist, but we only track and name the useful ones.

Your shoe plus the Eiffel Tower exists as a fusion, but it's not salient. It doesn't appear in our conceptual schemes because it serves no purpose. We don't

need to track scattered, non-functional composites. But the fusion is there in the ontology, ignored but real.

This preserves the theoretical benefits of unrestricted composition (no vagueness, no arbitrary cutoffs) while acknowledging common sense (we don't actually care about arbitrary fusions). The cost is admitting that reality is vastly more populated than our conceptual schemes recognize. Most objects are invisible to us, not because they don't exist, but because they're pragmatically irrelevant.

This middle ground is increasingly popular among philosophers. It lets you have clean metaphysics without pretending that human concepts carve nature at its joints. Nature has no joints. It has smooth continuity everywhere. We impose joints by selecting certain fusions as salient. The selection is mind-

dependent, but the fusions are not.

The downside is that this makes almost everything we say about objects pragmatic rather than metaphysical. When you say "there's a desk here," you're not reporting a fundamental feature of reality. You're highlighting one fusion among infinite overlapping fusions because that particular fusion matters to you. The desk is real, but so are desk-minus-one-atom, desk's left half, desk-plus-random-atoms-in-Paris, and infinite others.

If you find this conclusion disturbing, you're not alone. It means the world of everyday objects is radically different from the world as it fundamentally is. We live in a pragmatically-selected subset of an ontologically bloated reality. Most of what exists, we ignore. What we attend to, we mistake for the whole of existence.

Living With Arbitrary Objects

Suppose unrestricted composition is true and arbitrary objects exist. How should you think and act differently?

Basically, you shouldn't. The existence of Deskapoleon doesn't affect anything. It has no causal powers beyond what its parts have. It doesn't do anything. It's not located anywhere useful. Knowing it exists doesn't help you predict, explain, or control anything.

This is the key insight: metaphysical ontology and practical ontology come apart. What exists metaphysically is vast and promiscuous. What exists practically is whatever you need to track for your purposes. The two ontologies don't conflict. They're answers to different questions.

Metaphysically, you might accept that infinite overlapping fusions occupy your

desk's location. Practically, you ignore all of them except the desk-fusion because that's the one that serves your purposes. The metaphysics doesn't stop you from living normally. It just reveals that "living normally" involves ignoring most of what exists.

The philosophical payoff is recognizing that object-talk is pragmatic. When we individuate objects, we're not discovering pre-existing boundaries in nature. We're imposing boundaries based on what matters to us. Different purposes, different boundaries. The objects are real (as fusions of their parts), but the boundaries are conventional.

This dissolves many pseudo-problems. Is a cloud an object? Metaphysically, yes (it's a fusion of water droplets). Practically, depends on your purposes. For meteorology, yes. For inventory, no. The question "is it really an object?" is

confused. The real question is “should we treat it as an object for this purpose?”

The Upshot

Arbitrary objects are philosophically embarrassing but metaphysically harmless. They bloat the ontology without affecting how the world works. Whether you accept them depends on whether you prioritize ontological parsimony or theoretical simplicity.

Parsimony says minimize your ontological commitments. Don’t accept Deskapoleon because it does no work. Keep only objects that appear in your best theories or matter for your practices. This is intellectually satisfying but creates vagueness problems.

Simplicity says accept all fusions to avoid vagueness and arbitrariness. Let the ontology bloat if necessary to keep the theory clean. This is theoretically

satisfying but requires accepting infinite objects you'll never think about.

Most people's intuitions favor parsimony. But most formal theories favor simplicity. This is a real tension with no easy resolution. You can't have both maximal parsimony and maximal simplicity. You have to choose what you care about more.

The practical lesson is to stop worrying about what "really exists" and start asking what objects you need to track for your purposes. The desk exists. So does desk-minus-one-atom. Both are equally real as fusions. But you only care about one because only one is salient for furniture-using practices.

Salience is not the same as existence. Most of what exists isn't salient. Most of what's salient isn't fundamentally real (it's a pragmatic carving of reality for human purposes). These facts are

compatible once you stop conflating metaphysical and pragmatic questions.

Arbitrary objects teach you that reality doesn't come pre-divided into the objects you perceive. You do the dividing. The world is rich enough to support infinite divisions. Your job is to figure out which divisions serve your purposes, not to discover the one true division that nature prefers.

Nature doesn't prefer anything. It just is. You're the one carving it up. Do it well.

6. Scattered Objects and Distributed Wholes

The state of California is real. It has borders, a government, laws, and forty million people. It does things: it taxes, builds infrastructure, passes regulations. You can be inside it or outside it. It's as real as any object you'll encounter.

It's also scattered across 163,000 square miles with massive gaps where nothing remotely California-like happens. There are forests with no government presence, private properties the state doesn't control, and federal lands that belong to a different jurisdiction. The borders are often invisible lines drawn on maps. Step across the line from California into Nevada and the object you were inside suddenly ends, replaced by a different object with different laws and different authority.

Is California a single object? If so, it's not the kind of object you can hold in your hand or see all at once. It's distributed, discontinuous, and defined by institutional facts rather than physical continuity. Welcome to the world of scattered objects, where location stops mattering and composition gets weird in new ways.

Are Nations Real Objects?

Start with the obvious case: nations. The United States exists. You can point to it on a map, visit it, be a citizen of it. It has mass (the combined mass of its territory and everything on it). It has location (these specific coordinates). It has causal powers (it can declare war, print currency, enforce laws).

But what exactly is the United States? It's not just the landmass, because the landmass existed before 1776 and wasn't the United States. It's not just the

government, because the government is made of people and those people change. It's not just the citizens, because citizens can leave the country and remain citizens.

The United States is something like: a legal-political entity defined by a constitution, instantiated in a government, recognized by other nations, and sovereign over a particular territory. It's a scattered object whose parts include land, people, buildings, laws, and institutions, all held together by collective recognition and coordinated behavior.

This makes nations ontologically strange. They're real enough to do things. They sign treaties, wage wars, build infrastructure. They have causal powers that can't be reduced to individual actions. When "the United States" invades another country, no individual

American invaded it. The nation did, acting through military institutions that execute collective decisions.

But nations are also sustained by belief and agreement. If everyone stopped recognizing the United States as a legitimate authority, it would cease to exist. The land would remain, the people would remain, but the political entity would dissolve. This makes nations both real and conventional, both causally powerful and dependent on collective intentionality.

Most objects aren't like this. Your desk doesn't depend on anyone recognizing it. It's a desk whether you believe in it or not. But California does depend on recognition. It exists because people collectively treat it as existing and coordinate their behavior accordingly. Remove the collective treatment and you have land with people on it, not a state.

Is a Cloud a Thing or a Mood?

Clouds are even stranger than nations because they're physically distributed and constantly changing.

A cloud is a visible mass of water droplets suspended in air. It has approximate boundaries (where droplet density drops below visibility). It has location (up there). It has properties (fluffy, gray, threatening rain). It's clearly an object in some sense.

But which water droplets are part of the cloud? Droplets are constantly condensing into the cloud and evaporating out of it. The cloud moves, which means different droplets occupy the cloud-region at different times. The boundaries are fuzzy. Zoom in and you can't tell where cloud ends and non-cloud air begins. There's just a gradient of droplet density.

Is the cloud the same cloud from moment to moment? It changes shape continuously. It gains and loses droplets constantly. It can split into multiple clouds or merge with others. After an hour, it might not share any droplets with its earlier self. Same cloud or different cloud?

Meteorologists treat clouds as objects because it's useful. Clouds appear on weather radar, move in trackable ways, and produce predictable effects (rain, shade, lightning). The object-talk serves prediction and explanation. But the object is fuzzy in ways that make metaphysicians uncomfortable.

The cloud is real in that it's a physically sustained pattern of water droplets held aloft by atmospheric conditions. But it's not real in the way a rock is real. The rock has clear boundaries and stable composition. The cloud has fuzzy

boundaries and constantly changing composition. The rock persists through structural integrity. The cloud persists through pattern stability in a dynamic system.

This suggests a spectrum of objecthood. On one end, you have paradigm objects like rocks: well-bounded, stable, materially persistent. On the other end, you have patterns in dynamic media: clouds, waves, whirlpools, traffic jams. The patterns are real and causally relevant, but they're sustained by flow rather than structure. They exist as long as conditions maintain the pattern, and they dissolve when conditions change.

Digital Files, Playlists, Corporations, Gods

Now leave the physical domain entirely and look at objects that exist in abstract or institutional spaces.

A digital file is scattered across your hard drive in non-contiguous sectors. The operating system presents it as a unified object (“Document.docx”), but physically it’s fragments distributed across the disk, linked by a file allocation table. When you move the file, you’re not moving matter. You’re updating pointers in the file system. The file exists as an organizational layer on top of magnetic patterns on a platter.

Is the file a real object? It has identity (this specific file). It has properties (size, creation date, content). It has location (sort of, on this drive in these sectors). You can copy it, which seems to create a new instance of the same file. You can delete it, which destroys the file while leaving the matter untouched. The file is real, but it’s not made of matter. It’s made of information organized in particular ways.

A playlist is even more abstract. It's a list of references to songs, not the songs themselves. The same song can appear in multiple playlists without being duplicated. Rearrange the playlist and you haven't changed the songs, just their ordering. Delete the playlist and the songs remain. The playlist is a purely organizational object with no material instantiation beyond whatever data structure represents it.

Corporations are legally real objects that exist almost entirely through documentation and agreement. A corporation has assets, liabilities, obligations, and rights. It can own property, enter contracts, and be sued. But it has no physical location in any simple sense. Its headquarters is a building the corporation owns, not the corporation itself. Its employees are people who work for the corporation, not

parts of the corporation. The corporation is a legal entity that exists because laws say it exists and people behave as if it exists.

Gods, if they exist, are maximally scattered. Traditional theism says God is omnipresent, which means God is located everywhere simultaneously. This makes God a scattered object occupying all of space. Or maybe God transcends space entirely, which means God isn't located anywhere. Either way, God's relationship to space and parts is nothing like ordinary objects. God is either maximally scattered or not spatial at all.

These examples show that scattered and distributed wholes come in wildly different varieties. Some are physical (clouds), some are informational (files), some are institutional (corporations), some are possibly supernatural (gods). What they share is that standard

mereology, designed for paradigm physical objects, strains to describe them.

Why Location Stops Mattering Faster Than You Think

For ordinary objects, location matters enormously. The hammer is in the toolbox or it's not. The difference determines whether you can use it. Location is tied to identity: this hammer here is a different hammer than that hammer there.

But location becomes negotiable quickly once you leave paradigm cases.

Your bank account is located… where? At the bank's headquarters? On a server? Distributed across multiple servers in different locations? The question doesn't have a good answer because bank accounts aren't spatially located in any straightforward way. They exist in information systems, and those systems

are distributed for redundancy and efficiency.

Your reputation is located… where? Wherever people think about you? In social media databases? In the minds of everyone who knows you? Reputation is a social object that exists through distributed beliefs. It has no location except "in society," which is not a place.

The English language is located… where? Everywhere English is spoken? In the brains of English speakers? In written texts? In the abstract space of possible linguistic structures? English is real and causally relevant (it shapes how people communicate and think), but it's not located anywhere specific. It's instantiated in millions of places simultaneously, and none of those instantiations is the language itself.

Once you accept that objects can be scattered, distributed, or non-spatial,

location stops being a fundamental property. Some objects have location. Some have distributed location. Some have no clear location at all. Location is a feature of certain kinds of objects, not a requirement for objecthood.

This has implications. It means you can't use location to individuate all objects. You can't point to everything and say "that's where this object is." Some objects are where they're instantiated, which might be multiple places or no particular place. Identity and location come apart.

When Parts Are Nowhere Near Each Other

Scattered objects raise a question: do parts need to be near each other to compose a whole?

For physical wholes, proximity usually matters. The parts of your car are

spatially close and connected. Scatter them across a field and you don't have a car anymore, you have car parts. Proximity and connection are part of what makes physical wholes cohere.

But this is contingent, not necessary. Unrestricted composition says any collection of parts composes a whole, regardless of location. Your left shoe plus the Eiffel Tower is a scattered object with parts in different locations. The spatial separation doesn't prevent composition, it just makes the object useless for most purposes.

Nations have this feature. Alaska is part of the United States despite being separated from the continental U.S. by Canada. The separation doesn't prevent Alaska from being part of the nation because national composition isn't based on spatial contiguity. It's based on political authority and legal jurisdiction.

Digital systems have this feature. A distributed database has nodes in different data centers across the planet. The nodes compose a single database because they're logically connected and coordinated, not because they're spatially close. The database is one object with scattered implementation.

Teams have this feature. A remote-working company has employees scattered globally. They compose a company because they coordinate toward shared goals under shared management. The company is real (it signs contracts, generates revenue, has legal identity) despite its parts never being in the same room.

This suggests proximity is pragmatically important but metaphysically optional. For objects that depend on causal interaction, proximity matters because interaction gets harder at distance. But

for objects that depend on informational or institutional relationships, proximity is irrelevant. The relationships can span any distance as long as communication is possible.

The Role of Coordination Across Distance

What holds scattered objects together if not spatial proximity? Usually, coordination.

Flocking birds are scattered across the sky but compose a flock through behavioral coordination. Each bird matches velocity with neighbors, creating a unified motion pattern. The flock has properties (direction, shape, speed) that emerge from coordination rather than proximity.

Markets are scattered across the planet but compose a global economy through price coordination. Trades in Tokyo

affect prices in New York through information flow. The market is unified by coordination mechanisms (arbitrage, information systems, shared currencies) rather than physical connection.

The internet is maximally scattered but composes a single network through protocol coordination. Every device running TCP/IP can communicate with every other such device. The network is unified by shared communication standards, not by proximity or physical connection.

Coordination is a compositional mechanism that works at any distance. It requires information flow rather than physical contact. As long as parts can exchange information and adjust behavior accordingly, they can compose a coordinated whole even when scattered across arbitrary distances.

This makes coordination-based wholes

fragile in specific ways. Break the communication channels and the whole dissolves. The flock scatters when birds can't see neighbors. The market fragments when information can't flow. The internet partitions when cables are cut. The whole is only as stable as the coordination mechanisms maintaining it.

Institutions as Distributed Wholes Held Together by Agreement

The strangest scattered objects are institutions. They have no physical location, no material composition, and no continuous existence independent of people recognizing them. Yet they're as real as anything in social life.

Marriage is an institution. It exists across all cultures that practice it. It's not located anywhere. It's not made of matter. It exists as a social practice sustained by collective recognition and behavior. Remove the collective

recognition and marriage disappears, even though the people involved and the ceremonies performed continue to exist.

Money is an institution. Dollar bills are physical, but money is not. Money is whatever a community treats as a medium of exchange and store of value. Shells have been money. Cigarettes in prison are money. Bitcoin is money in some contexts. The physical substrate varies, but the institution persists as long as people coordinate their treatment of the substrate.

Universities are institutions. Yale exists, but where? In the buildings? Students and faculty leave and enter constantly. The buildings could burn down and Yale would rebuild. In the endowment? That's just money, and money is itself institutional. Yale exists as a set of practices, relationships, and authorities sustained by ongoing recognition and

coordination.

These institutions are real because they constrain behavior and enable action. You can't just declare yourself married and have it be legally recognized. You can't print your own money and force stores to accept it. You can't award yourself a Yale degree. The institutions have authority independent of any individual's will, even though they're sustained by collective will.

This makes institutional objects ontologically peculiar. They're causally powerful but non-physical. They're real but dependent on belief. They're objective (you can't change them by yourself) but socially constructed (they'd disappear if everyone stopped believing in them together). They exist in a space between physical reality and pure fiction.

The Uncomfortable Possibility That Most Objects Are Like This

Here's the disturbing thought: maybe all object boundaries are like institutional boundaries. Maybe the only difference between nations and rocks is that national boundaries are obviously conventional while rock boundaries seem natural.

But are rock boundaries natural? Zoom in and you find atoms. Atoms are held together by forces, but forces are continuous fields. There's no sharp boundary where rock ends and air begins, just a gradient of atomic density. The "rock" is a pattern of matter that we treat as a unified object because it's causally cohesive and perceptually salient.

You might object: the rock would be there even if humans didn't exist. True. The matter would be there, arranged in that configuration. But would there be a rock-object, distinct from the matter? Or would there just be matter, arranged

rock-wise, which we carve into an object called “rock”?

This is the deep question scattered objects raise. If scattered objects exist only through coordination, agreement, or collective recognition, and if those mechanisms are obviously mind-dependent, then scattered objects are mind-dependent. But if the only difference between scattered and non-scattered objects is proximity, and proximity is just a pragmatic feature, then maybe all objects are mind-dependent in the same way.

Not mind-dependent in their material substrate. The matter is there regardless. Mind-dependent in their boundaries, unity, and identity. We carve the matter into objects based on what matters to us. Some carvings correspond to natural joints (stable patterns, causal cohesion, functional integration). But the decision

to treat these patterns as unified objects rather than mere patterns is ours.

Living in a World Without Clear Boundaries

Once you see that scattered objects are real, you can't unsee it. And once you realize they're everywhere, the world looks different.

Your self is scattered. Your body is here, but your reputation is distributed across social networks. Your memories are partly in your brain and partly in photo albums and digital storage. Your identity is partly constituted by relationships with people who aren't spatially near you. You are not a bounded, local object. You're a distributed pattern spanning physical, social, and informational spaces.

Your projects are scattered. You're writing a book, but the book exists partly

in files, partly in notes, partly in your head, partly in conversations with editors. The book is real before it's published, but it's scattered across multiple substrates. Completing the book means gathering these scattered parts into a unified form.

Your problems are scattered. Climate change is a distributed whole with causes and effects spanning the planet. No one location contains the problem. No one action solves it. The problem exists as a global pattern that can only be addressed through distributed coordination.

Recognizing scattered objects means recognizing that most things worth caring about aren't neatly bounded. They leak across spatial, temporal, and categorical boundaries. They exist partly here and partly elsewhere. They're held together by relationships, coordination, and agreement rather than physical

connection.

This is uncomfortable because we want clean categories and clear boundaries. We want to point at things and say “this is one object, that is another.” But the world doesn’t respect our desire for cleanliness. Objects blend into each other. Wholes contain wholes contain wholes, all with different boundaries serving different purposes.

The upshot is to stop expecting objects to be spatially unified. Some are. Many aren’t. The scattered ones are no less real. They’re just real in ways that require different conceptual tools to track and understand. Tools that acknowledge distribution, coordination, and convention as legitimate compositional mechanisms alongside proximity and bonding.

Location stops mattering when you realize composition doesn’t require it.

What matters is the right kind of relationship between parts, and that relationship can span any distance. The world is full of scattered, distributed, coordinated wholes doing real work despite having no single location. Once you see them, you realize they're not edge cases. They're everywhere, hiding in plain sight because we've been looking for the wrong kind of unity.

7. Biology: Organisms as Negotiated Wholes

You are not one thing. You are approximately 37 trillion human cells, plus another 38 trillion bacterial cells, plus viruses, fungi, archaea, and other passengers whose numbers nobody has accurately counted. You are a walking ecosystem that has convinced itself it's a person.

This should disturb you more than it does. The boundaries of your body are not where you think they are. The things you consider "you" are outnumbered by things you consider "not you," and the distinction between the two is increasingly difficult to defend. Your survival depends on entities you're hosting, and their survival depends on you. The relationship is so intimate that calling it symbiosis undersells the

ontological confusion.

Biology is where mereology gets messy in ways philosophy hasn't fully processed. Organisms aren't tidy wholes with clear boundaries and stable parts. They're negotiated compromises between competing entities, held together by evolutionary deals that could collapse at any moment. This chapter is about what actually composes your body and why the answer is more disturbing than "cells."

Cells, Organs, Symbionts, Microbiomes

Start with the standard story: you are composed of cells, cells compose into tissues, tissues compose into organs, organs compose into organ systems, organ systems compose into you. Clean hierarchy. Neat levels. Each level composed of units from the level below.

This story is approximately correct and deeply misleading. The hierarchy exists, but it's riddled with exceptions, ambiguities, and entities that don't fit cleanly anywhere.

Red blood cells have no nucleus and can't reproduce. Are they cells or cell fragments? They're alive in some sense (metabolically active), dead in others (can't divide, will die in ~120 days). They're produced by your bone marrow and destroyed by your spleen, circulating through your body while doing exactly one job: carrying oxygen. They're you, clearly, but they're also more like disposable tools your body manufactures than like the cells that compose tissues.

Cancer cells are you genetically (same DNA) but they're rebelling against you functionally. They grow without regard for organismal integrity. They metastasize, which is a polite term for

colonizing your body like invaders. They evolve to evade your immune system. At what point do they stop being “you” and become a separate organism that happens to share your genome? The question isn’t academic. It affects how we think about cancer treatment: are we repairing you or fighting an invasion?

Gut bacteria are not you genetically (different DNA, different evolutionary lineage), but they’re essential for your survival. They digest food you can’t digest, produce vitamins you can’t synthesize, train your immune system, and influence your mood through biochemical signals to your brain. Remove them and you get sick. They’re as essential to your functioning as your liver, but nobody thinks bacteria are organs.

Where does this leave the boundary between you and not-you? The standard

answer is genetic: same DNA means you, different DNA means not-you. But this doesn't work. Gut bacteria are different DNA, essential for survival. Mitochondria are different DNA (they have their own small genome), essential for survival, and universally considered part of your cells. Cancer is same DNA, trying to kill you. Skin cells you shed constantly are same DNA, clearly not you anymore once they're floating in the air.

The boundary is negotiated, not given. It depends on what you're trying to explain and what criteria matter for your purposes.

When Does a Body End?

The skin seems like an obvious boundary. Inside the skin is you, outside is not-you. Simple.

Except your skin is covered with

bacteria. The skin microbiome contains millions of organisms living on your surface, in your pores, feeding on dead skin cells and oils. They're not inside you, but they're on you, and they're part of what makes your skin function. Some protect against pathogens. Others produce compounds that affect skin health. Remove them (say, with antibiotics) and your skin becomes vulnerable.

Are these bacteria part of your body? They're outside the skin barrier, so technically external. But they're so closely associated with you that they're more like an additional layer than like independent organisms. They travel with you. They depend on you. You depend on them. The boundary blurs.

Now consider your gut. The lining of your intestines is inside your body in a topological sense, but it's exposed to the

external environment. Food you eat enters your mouth (external environment), passes through your digestive tract (still topologically external, like water through a hose), and exits without ever crossing into your internal tissues. Your gut is a tunnel through your body, not truly inside it.

The bacteria in your gut are in that tunnel. They're topologically external but functionally essential. They're processing food that you ate but can't digest. They're producing metabolites that cross into your bloodstream and affect your biology. Are they inside you or outside you? The topology says outside. The function says inside. There's no fact of the matter.

Then there's the placenta during pregnancy. The placenta is partly fetal tissue (genetically the baby, not the mother) and partly maternal tissue. It sits

at the interface between mother and fetus, mediating exchange of nutrients and waste. The fetus depends on it. The mother's body built it. Who does it belong to? Both? Neither? It's a shared organ created by two organisms for mutual benefit, lasting only as long as the pregnancy.

Your body doesn't end at your skin. It ends at a fuzzy, negotiated boundary that depends on what you count as functional integration versus mere association. Some parts are clearly inside (liver, brain). Some are clearly outside (the chair you're sitting on). Vast numbers are in between: on you, in you topologically-externally, exchanged with you, produced by you but now separate.

Cancer as a Mereological Rebellion

Cancer clarifies the problem because it's you rebelling against you.

Cancer cells start as your cells. They accumulate mutations that disable growth regulation. Instead of dividing when needed and stopping when done, they divide continuously. Instead of staying in their tissue, they invade surrounding tissues. Instead of dying when damaged, they evade apoptosis (programmed cell death).

From the cancer cell's perspective, it's optimizing for its own reproduction. Your body is just an environment to exploit. The cancer isn't trying to kill you, it's just indifferent to your survival. You die as a side effect of the cancer's success.

From your perspective, the cancer is a traitor. It's breaking the cooperative deal that makes multicellular organisms possible. That deal is: cells sacrifice their individual reproductive interests for the good of the organism, and in exchange

they get resource support and protection. Cancer cells defect from this deal. They take resources without contributing to organismal function.

This makes cancer a mereological rebellion. The parts are defecting from the whole. They're optimizing locally (cell-level reproduction) at the expense of global function (organism survival). The whole is composed of parts, but the parts have their own interests, and those interests sometimes conflict with the whole's interests.

This shouldn't surprise you. Every cell in your body is descended from a lineage of single-celled organisms that survived by reproducing. Multicellularity is a recent evolutionary innovation (last 600 million years or so). Your cells retain the machinery for independent reproduction. They're just normally suppressed by regulatory systems. Cancer happens

when suppression fails.

You are not a unified agent. You are a temporarily aligned collective of agents (cells) that mostly cooperate because cooperation was fitness-enhancing for their ancestors. When cooperation breaks down (cancer), the collective fractures. Some parts pursue their own interests at the expense of the whole. The whole disintegrates unless external intervention (medicine) restores cooperation or eliminates defectors.

This is true of all multicellular organisms. We're not naturally unified. We're evolutionarily coerced into cooperation by mechanisms that suppress part-level interests in favor of whole-level interests. When those mechanisms fail, the organism comes apart.

Evolution as Part-Level Opportunism

Evolution doesn't care about wholes. It

cares about reproduction, which happens at the level of genes and sometimes organisms. This creates constant tension between part-level and whole-level selection.

Mitochondria are the clearest example. They used to be independent bacteria. About 1.5 billion years ago, one bacterium engulfed another (probably trying to eat it). The engulfed bacterium survived inside and started providing energy through aerobic respiration. The host provided protection and resources. Both benefited. The relationship became permanent.

Now mitochondria can't survive outside cells, and eukaryotic cells (the kind you're made of) can't survive without mitochondria. What was two organisms became one organism with organelles. The boundary between organisms and parts dissolved through evolutionary

integration.

But mitochondria retain their own genome, separate from yours. They reproduce independently inside your cells. They're transmitted only through the maternal line (sperm don't contribute mitochondria). This creates conflicts. Mitochondrial genes are selected to maximize maternal-line transmission, which isn't always aligned with organismal fitness.

Some mitochondrial mutations cause disease. Your nuclear genes (in the cell nucleus) would be selected to suppress these mutations, but they can't because mitochondria control their own reproduction. The result is tension between genomic factions within your cells. You're not genetically unified even at the cellular level.

This pattern repeats everywhere in biology. Organisms are coalitions of

genes, and genes are selected for their own replication, not for organismal benefit. Most of the time these interests align. Sometimes they don't, and you get conflicts: parent-offspring conflict, sexual conflict, intragenomic conflict.

Your genome contains about 20,000 genes. They cooperate most of the time because they're inherited together (you reproduce as a whole, transmitting the whole genome). But they're not perfectly aligned. Some genes are selected to spread even at the cost of organismal fitness. Transposable elements ("jumping genes") copy themselves around the genome, creating mutations. The organism would be better off without them, but they're hard to eliminate because they're good at copying.

Evolution is not designing organisms as unified wholes. It's selecting genes for replication, and organisms are the

temporary vehicles genes build to achieve replication. The vehicles have to be functional (dead organisms don't reproduce), so genes are selected to build coherent organisms. But the coherence is a side effect of selection, not the goal.

This means organisms are always vulnerable to part-level exploitation. Whenever a gene can increase its replication at the organism's expense without getting purged by selection, it will. Organisms are held together by evolutionary mechanisms that suppress cheating, not by natural unity.

When Parts Can't Be Excused by the Whole

The fact that organisms are negotiated coalitions has moral implications. You can't excuse part-level behavior by appealing to whole-level necessity if the whole is just a coalition of parts pursuing their interests.

When cancer cells defect and kill the organism, we don't excuse them by saying "they were just following their evolutionary interests." We treat them as pathological, something to eliminate. The cancer has reasons (maximize reproduction), but the reasons don't justify the behavior.

This asymmetry is interesting. We treat the organism as the morally relevant whole and cells as parts that should subordinate to the whole. When parts defect (cancer), we side with the whole and destroy the parts. We don't say "cancer cells have rights too" or "from their perspective, they're just trying to survive."

But organisms are also parts of larger wholes: populations, ecosystems, societies. When an organism defects from ecosystem function (say, an invasive species), we sometimes think it

should be controlled or eliminated for the good of the ecosystem. The organism has interests, but those interests are subordinate to ecosystem health.

This creates a ladder of wholes, each containing parts that have their own interests. Cells compose organisms, organisms compose populations, populations compose ecosystems. At each level, there's tension between part-level interests and whole-level interests. And at each level, we have to decide which level's interests take priority.

There's no objective answer. Different contexts prioritize different levels. Medicine prioritizes the organism. Conservation prioritizes the ecosystem. Evolution doesn't prioritize anything, it just selects what reproduces.

The mereological lesson is that you can't automatically privilege the whole over the parts or the parts over the whole.

Both are real. Both have interests. Which you prioritize depends on your values and purposes, not on metaphysical facts about which level is "more real."

Symbiosis as Compositional Confusion

Symbiosis is the term biologists use for close relationships between different species. It comes in three flavors: mutualism (both benefit), commensalism (one benefits, other unaffected), and parasitism (one benefits, other harmed).

But symbiosis creates ontological confusion. When organisms are so tightly integrated that they can't survive separately, do they compose a single organism or remain separate organisms in a relationship?

Lichens are the classic case. A lichen is a fungus and an algae (or cyanobacterium) living together. The fungus provides structure and protection. The algae

photosynthesize and provide energy. Neither can survive in the environments where lichens live without the other. They reproduce together (lichen fragments contain both partners). They grow together as a unified structure.

Is a lichen one organism or two? Biologists can't agree. Genetically, it's two organisms (different DNA). Functionally, it's one organism (integrated metabolism, unified reproduction, shared fate). The question doesn't have a clean answer because the categories don't fit.

Coral is similar. What looks like a coral colony is actually thousands of individual polyps (the coral animals) plus millions of algae (zooxanthellae) living inside their tissues. The polyps build the structure. The algae photosynthesize and provide most of the coral's energy. Both depend completely on each other. The

coral bleaches (turns white and dies) when the algae leave.

Are corals one organism or many? Each polyp is genetically identical to other polyps in the colony (they're clones). But they're separate animals, connected by tissue but individually capable of feeding and reproduction. Plus the algae, which are genetically distinct but functionally integrated. You have a colonial animal hosting millions of algal symbionts, all functioning as an ecological unit. One organism? Many? A superorganism? The language struggles.

Humans are milder versions of the same phenomenon. We're not as dependent on our microbiomes as lichens are on their partners, but the dependence is real. Germ-free mice (raised without any bacteria) are viable but abnormal. Their immune systems don't develop properly. Their guts are malformed. They're

missing metabolic capabilities. They're alive but not fully functional.

Are you one organism (human) or a holobiont (human plus microbiome functioning as a unit)? For most purposes, saying "I am an organism" works fine. But for understanding your actual biological function, you need to think of yourself as a community. The community has a primary member (you, genetically), but it has essential supporting members (microbiome) without which the primary member fails.

Biology keeps discovering that what looked like clean organismal boundaries are actually porous, negotiated interfaces where multiple entities blend together. The more closely we look, the less clear "organism" becomes as a category.

The Uncomfortable Truth About Biological Individuality

Here's what biology has learned: biological individuality is a spectrum, not a binary. Some entities are clearly individual organisms (you, mostly). Some are clearly non-individuals (populations of unrelated bacteria). Many are in between.

A Portuguese Man O' War looks like a jellyfish but is actually a colonial organism: multiple polyps (genetically identical) specialized for different functions (feeding, reproduction, defense, flotation). Each polyp is an individual in some sense (it's a separate animal). But the colony is also an individual (it functions as a unit, dies as a unit, can't be decomposed into viable polyps).

Aspen groves look like forests of individual trees but are often a single organism. The trees are genetically identical stems connected by

underground root systems. The whole grove is one genetic individual with thousands of trunks. It reproduces clonally (new stems sprout from roots) and can live thousands of years. Is each trunk an individual? Or is the whole grove the individual?

Slime molds spend part of their life as individual amoebae and part as a multicellular slug. When food is scarce, individual amoebae aggregate into a slug that moves toward light and forms a fruiting body. Some amoebae become the stalk (and die). Some become spores (and reproduce). During the multicellular phase, are they one organism or many? They're cooperating genetically unrelated individuals forming a temporary multicellular structure. Then they disperse and become individuals again.

These examples show that individuality

is not a fundamental biological category. It's a pragmatic one we impose based on functional integration, genetic relatedness, and reproductive independence. Different criteria give different answers about what counts as an individual.

The unsettling implication is that you are an individual by convention and context, not by metaphysical necessity. For most purposes (medical, legal, social), treating you as an individual makes sense. But biologically, you're a nested hierarchy of entities with partial autonomy at every level. Genes pursue their interests. Cells pursue theirs. Organs coordinate for organismal function. The organism coordinates with microbiomes for holobiont function.

There's no level that's "the real individual." Each level is real and has its own dynamics. Which level matters

depends on what you're trying to explain. Evolution operates on genes. Medicine operates on organisms. Ecology operates on populations. Each level is the right level for its domain.

The Takeaway

Biology teaches that organisms are not natural kinds with essences. They're evolved coalitions held together by mechanisms that suppress part-level exploitation. The coalitions are stable enough to persist for an organism's lifetime (usually), but they're always vulnerable to part-level defection (cancer, senescence, intragenomic conflict).

The boundaries of organisms are negotiated between what counts as functionally integrated versus merely associated. Some parts are clearly inside (brain, liver). Some are clearly outside (soil bacteria). Many are in between

(microbiome, mitochondria, viruses integrated into your genome).

This makes “you” a fuzzy, context-dependent category. Genetically, you’re the cells with your nuclear DNA. Functionally, you’re the human cells plus microbiome plus mitochondria. Topologically, you’re whatever is enclosed by a continuous boundary, which excludes gut bacteria. Evolutionarily, you’re a vehicle for gene reproduction, not an individual agent.

All these boundaries are real. None is uniquely correct. Which boundary matters depends on your explanatory purposes. Medicine wants genetic boundaries (to target cancer specifically). Nutrition wants functional boundaries (because microbiome affects digestion). Evolution doesn’t care about organisms, only genes and populations.

The philosophical upshot is that

biological wholes are negotiated achievements, not given facts. Organisms maintain themselves through active regulation of part-level behavior. When regulation fails, the organism disintegrates into parts pursuing their own interests. The whole exists only as long as the parts can be kept in line.

This is true for all evolved coalitions, biological or social. You're not naturally a unified whole. You're a temporary alliance of entities that mostly cooperate because cooperation was selected for. When cooperation breaks down, the whole dissolves. Wholeness is something achieved, maintained, and defended, not something given. Biology shows you the mechanisms by which that happens, and the ways it fails.

8. Technology: Modular Reality

Your phone is designed to break in exactly the right way. Not break completely, which would be bad for Apple's reputation. Break modularly, which is perfect for Apple's business model. The screen shatters but the processor survives. The battery degrades but the camera still works. The charging port fails but everything else functions. Each failure is an isolated event that can be repaired (or more likely, that motivates replacement).

This is not an accident. Modularity is the fundamental organizational principle of modern technology. Complex systems are built from components that can be designed, manufactured, tested, replaced, and upgraded independently. The components interact through defined interfaces, and as long as the interfaces

are respected, you can swap components freely without understanding how they work internally.

Modularity is mereology made practical. It's a theory of how parts compose into wholes when the wholes need to be buildable, maintainable, and extensible by humans with bounded intelligence. This chapter is about what modularity reveals about composition and why systems break at the seams, not the center.

Machines as Replaceable Parts With Personalities

Start with the obvious: machines are assemblies of parts. You can disassemble a car into components, fix or replace the broken one, and reassemble it. The car's identity persists through part replacement because the organization persists.

But machines are not just assemblies.

They’re hierarchical assemblies with personalities at every level. A carburetor is not just “a part of an engine.” It’s a subsystem with its own internal complexity, its own failure modes, and its own repair procedures. The carburetor has jets, needles, floats, and chambers, each of which is a part that can fail independently.

This creates a parts hierarchy: car composed of systems (engine, transmission, brakes), systems composed of subsystems (carburetor, distributor, fuel pump), subsystems composed of components (jets, springs, gaskets), components composed of materials (steel, rubber, brass). At each level, the parts have functional roles and interact through defined interfaces.

The genius of modularity is that you don’t need to understand the whole hierarchy to work at one level. A

mechanic replacing a carburetor doesn't need to understand metalworking or combustion chemistry. They just need to know: remove old carburetor, install new carburetor, ensure connections are correct. The carburetor is a black box with specified inputs and outputs. What happens inside is the carburetor's problem.

This is possible because the interfaces are standardized. The fuel line connects to the carburetor in a specific way. The air intake has a defined size. The mounting bolts follow a pattern. As long as the replacement carburetor matches these specifications, it will work. The internal design can be completely different (newer technology, different manufacturer, different materials), and it doesn't matter. Interface compatibility is all that's required.

But here's what's strange: the carburetor

has a personality. Two carburetors meeting the same specifications can behave differently. One might idle rough. Another might flood easily. A third might be finicky about adjustment. They're functionally equivalent (they all mix air and fuel), but they're not identical. The personality comes from tolerances, manufacturing variations, and emergent properties of the specific configuration.

This matters because it means modularity is approximate, not perfect. You can swap modules and mostly get equivalent behavior. But "mostly" is doing a lot of work. The replacement carburetor might require adjustment. It might have different quirks. The car will behave slightly differently. True interchangeability is an ideal that reality approximates without fully achieving.

Software as Pure Mereology

Software is mereology without the

compromises of physical reality. A program is composed of modules (functions, classes, libraries), modules are composed of statements, statements are composed of expressions, expressions are composed of operators and operands. Pure composition all the way down to bits.

Software modularity is spectacularly successful because the interfaces are perfect. A function doesn't care who calls it or what hardware it runs on. It takes inputs, produces outputs, and the internal logic is completely hidden. Call the same function a million times and you get identical behavior (assuming no hidden state or randomness). This is impossible with physical parts.

The cost of perfect interfaces is that software breaks in distinctive ways. Physical systems degrade gradually. Metal fatigues, rubber cracks, lubricants

evaporate. You get warning signs before total failure. Software doesn't degrade. It works perfectly until it encounters a case the programmer didn't anticipate, and then it fails catastrophically.

This is because software composition is logical rather than physical. A physical system can tolerate small deviations from specification. A bolt can be slightly too small and still work, just with more play. Software cannot tolerate deviations. An off-by-one error (wrong boundary condition in a loop) will cause incorrect behavior every time, with no warning and no graceful degradation.

Software also composes at scales that are psychologically impossible. A modern web application might depend on hundreds of libraries, each of which depends on other libraries. The dependency tree is thousands of nodes deep. No human understands the whole

system. The developers understand their code and the APIs of the libraries they directly use. They trust that the libraries work, and they trust that dependencies are managed correctly.

This creates a vulnerability: the dependency is real but the understanding is not. Your application depends on a library for parsing dates. That library depends on another library for time zone handling. That library depends on a third library for locale detection. Somewhere deep in the tree, a library has a bug or a security vulnerability. The bug propagates up through dependencies, affecting your application in ways you can't predict because you don't understand the internal behavior of libraries twelve layers down.

This is compositional fragility. The parts work independently, but the composition introduces emergent failures that appear

only when specific parts interact in specific ways. The failures are deterministic (same inputs always produce same outputs), but they're unpredictable because the interaction space is too large to fully explore.

APIs, Plug-ins, and Compositional Failure

Application Programming Interfaces (APIs) are formalized contracts between software components. The API specifies: here's what you can ask me to do, here's the format of requests, here's what I'll return, here's what errors I might signal. Obey the contract and we can work together without knowing each other's implementation.

APIs make composition safe-ish. You can write code that uses a database without knowing how the database stores data. You can use a payment processor without understanding credit card

networks. You can call a mapping service without knowing routing algorithms. The API hides complexity behind a simple interface.

But APIs are promises, and promises can be broken. An API might specify that a function returns a number, but fail to specify what happens when the input is negative. The caller assumes negative inputs are handled (why wouldn't they be?). The implementation assumes negative inputs won't occur (they're not mentioned in the spec). Both assumptions are reasonable and incompatible. The system works until someone passes a negative number, and then it fails in a way that surprises everyone.

This is compositional failure at the interface. The parts work individually. The interface specification is incomplete. The incompleteness creates a gap where

assumptions on each side don't match. The failure appears only when the gap is exposed.

Plug-ins are even more fragile. A plug-in is code that extends an application, written by someone other than the application developers. The plug-in uses APIs exposed by the application. The application has no idea what the plug-in will do, so it can't anticipate failure modes.

This creates adversarial composition. The application tries to protect itself by limiting what plug-ins can access. The plug-in tries to be useful by doing things the application developers didn't anticipate. The tension produces endless problems: plug-ins that crash the application, security vulnerabilities from plug-ins accessing sensitive data, performance degradation from poorly-written plug-ins.

The standard solution is sandboxing: run plug-ins in a restricted environment where they can't damage the application. But sandboxing limits what plug-ins can do, which limits their usefulness. The more powerful you make plug-ins, the more vulnerable the application becomes. The safer you make the application, the less powerful plug-ins can be. There's no perfect balance.

Web browsers face this constantly. Extensions (plug-ins for browsers) can modify any webpage, intercept network traffic, and inject code. This power makes them useful (ad blockers, password managers, accessibility tools). It also makes them dangerous (they can steal passwords, inject malware, track your browsing). Browsers try to sandbox extensions, but extensions keep finding ways around the sandbox because they need power to be useful.

This is the fundamental problem of compositional security: parts need access to compose, but access creates vulnerability. You can't have both perfect security and perfect composability. You pick a point on the spectrum and accept the tradeoffs.

Why Systems Break at the Seams, Not the Center

Well-designed systems have stable, heavily-tested cores and fragile, constantly-changing peripheries. The core components are mature, their failure modes are known, and they're protected by layers of abstraction. The periphery is where new features are added, where different subsystems interact, and where assumptions meet reality.

This is why failures cluster at interfaces. The core of a database might be rock-solid. The integration between the database and the web server is where

things break. Not because the integration is poorly designed, but because it's where two different systems with different assumptions have to coordinate.

Each system was designed independently, with different priorities. The database prioritizes consistency and durability. The web server prioritizes low latency and high throughput. These priorities are in tension. The integration has to reconcile them, which means making tradeoffs neither system was designed for.

The result is that the integration layer is brittle. It works under expected conditions. It fails under unexpected combinations of load, timing, and state. The failures are hard to debug because they involve both systems simultaneously. The database team says "our queries are fine." The web server team says "our code is fine." Both are

right. The integration is the problem.

This pattern repeats everywhere. Failures concentrate at module boundaries because boundaries are where different assumptions collide. Inside a module, one team made consistent decisions. Across modules, different teams made different decisions that seemed locally reasonable but are globally incompatible.

The car stalls at idle but runs fine at speed. The problem is the interaction between the idle control system and the fuel injection system, both of which work correctly in isolation. The web application crashes under high load, but only when users perform a specific sequence of actions. The problem is the interaction between session management and database transactions, both of which handle their individual concerns correctly.

Compositional failures are interaction

failures. The parts work. The composition doesn't, because the parts were never designed to work together in exactly this way. The integration assumes things the parts don't guarantee, or the parts assume things the integration doesn't provide. The gap between assumptions is where failures hide.

Abstraction Layers as Compositional Denial

Modern technology is built on abstraction layers: levels of organization where lower-level details are hidden and higher-level interfaces are exposed. You write code in Python, which compiles to bytecode, which runs on a virtual machine, which executes machine instructions, which control transistors. Each layer hides the layer below.

This is enormously productive. You can write Python without understanding virtual machines. You can write machine

code without understanding semiconductor physics. Each layer provides a stable interface that lets the layer above ignore the layer below's complexity.

But abstraction is compositional denial. You're pretending the lower layers don't exist or don't matter. Most of the time this works. Sometimes it fails catastrophically.

Your Python program runs slowly. You profile it and find a bottleneck. You optimize the Python code. It's still slow. The problem is that Python's virtual machine is doing something inefficient with your data structure. You can't fix it without understanding the virtual machine. The abstraction has leaked.

Or: your code works on your machine but fails in production. The difference is that production uses different hardware with different timing characteristics.

Your code made assumptions about operation order that happen to be true on your machine but aren't guaranteed by the abstraction. The hardware details you thought you could ignore turn out to matter.

Abstraction leakage is a mereological problem. You treated the system as composed of cleanly layered parts with perfect interfaces. Actually it's composed of loosely layered parts with approximate interfaces. The approximation is good enough most of the time. When it's not, you have to violate the abstraction and reach down to lower layers to fix problems.

This creates a dilemma. If you don't use abstraction, you're overwhelmed by complexity. If you do use abstraction, you're vulnerable to leakage. The middle path is using abstraction while remaining aware it's incomplete, and being ready to

pierce it when necessary. But this requires understanding multiple layers simultaneously, which is exactly what abstraction was supposed to prevent.

The best programmers maintain a mental model of multiple abstraction layers and know when to think at which layer. The code is at the Python layer. The performance problem is at the virtual machine layer. The deployment issue is at the operating system layer. The intermittent crash is at the hardware layer. You debug by moving up and down the stack, finding the right layer for each problem.

This is cognitively expensive. It's also unavoidable for any sufficiently complex system. Abstraction makes complexity manageable. It doesn't eliminate complexity. The complexity is still there, hidden but present, waiting to become relevant when the abstraction leaks.

Modularity as a Design Bet Against the Future

When you design a modular system, you're making a bet about where the boundaries will be stable. You're saying: these components will need to change independently, so I'll isolate them behind interfaces. These other components will always change together, so I'll couple them tightly.

If you guess right, modularity pays off. You can upgrade the database without touching the application. You can swap payment processors without rewriting checkout logic. You can replace the recommendation engine without affecting the user interface.

If you guess wrong, modularity creates overhead without benefits. You put an interface between two components that always change together anyway. Now you have interface maintenance costs

(versioning, documentation, compatibility) without the benefit of independent evolution.

Worse, if the world changes in ways that cut across your module boundaries, you have to refactor the entire system. The boundaries you thought were stable turn out to be artificial. The real boundaries are elsewhere. You have to decompose the system differently and recompose it along new lines.

This happens constantly. You design a system where user management is separate from content management. Then you need to integrate user preferences into content recommendations. The separation you created is now a barrier. You either violate the separation (tight coupling across the boundary) or refactor (move the boundary or create a new integration layer).

Modularity is a bet that certain

boundaries will remain stable and certain capabilities will remain independent. This bet is about future evolution, which is fundamentally unpredictable. You make your best guess based on current requirements and anticipated changes. The guess is often wrong, not because you're bad at guessing, but because the future doesn't respect your categories.

The systems that survive are the ones with boundaries that happened to match how the world evolved. The iPhone succeeded partly because its modular architecture (apps as independent components) matched how mobile software actually evolved. Other phones with different boundaries (tighter integration between system and applications) couldn't adapt as quickly.

This is selection on system architecture. Architectures that anticipate future evolution survive. Architectures that bet

on the wrong boundaries die or require expensive rewrites. You can't know in advance which architectures will win. You make your best guess and see what survives contact with reality.

When Composition Enables, When It Constrains

Modularity both enables and constrains. It enables by letting you work on parts independently. It constrains by forcing you to interact through interfaces.

Enabling: You're building a website. You split it into frontend (user interface) and backend (business logic and data). Two teams can work in parallel. Frontend team tries different designs without touching backend code. Backend team optimizes database queries without affecting frontend. The separation enables parallel development.

Constraining: Now you want to add a

feature that requires tight coordination between frontend and backend. Real-time updates, maybe. Every user action needs immediate backend feedback. Every backend change needs immediate frontend reflection. The separation you created now forces you to communicate through an API that adds latency and complexity. The feature would be easier to implement with tight coupling.

This is the modularity tradeoff. Loose coupling enables independent evolution but makes tight integration harder. Tight coupling makes integration trivial but prevents independent evolution. You can't have both. You pick the coupling that matches your expected evolution pattern and accept that some features will be harder than they would be with different coupling.

As systems grow, they accumulate features that cut across module

boundaries. Each feature is awkward to implement because it violates the original decomposition. Eventually you have more cross-cutting concerns than orthogonal concerns. The modularity that once enabled development now constrains it. You need to refactor, which means rebuilding the system with new boundaries that better match current concerns.

This cycle is inevitable. All successful systems eventually need architectural overhaul because they outlive their original design assumptions. The assumptions were reasonable when made. The world changed. The assumptions stopped matching reality. The architecture becomes a straightjacket.

The companies that survive are the ones that recognize when architecture has become constraint rather than enabler

and are willing to pay the cost of refactoring. The companies that die are the ones that cling to obsolete architectures because refactoring is expensive and risky.

The Upshot

Technology is applied mereology. It's about decomposing problems into parts, solving the parts independently, and composing solutions into wholes that work. The decomposition is never perfect. The interfaces leak. The assumptions fail. The boundaries become obsolete.

But modularity is still the best tool we have for managing complexity. It's the only known way to build systems too complex for any individual to understand. The alternative is unmanageable monoliths that collapse under their own weight.

The lesson is to treat boundaries as provisional. They're useful abstractions, not metaphysical truths. Good boundaries enable current work and anticipated evolution. When the world changes, boundaries need to change too. Clinging to boundaries because they used to work is how systems ossify and die.

This applies beyond technology. Organizations are modular (departments, teams, roles). Societies are modular (institutions, laws, markets). Scientific theories are modular (disciplines, subfields, specializations). All these boundaries are useful and artificial, enabling and constraining, temporary and essential.

Modularity is how we manage complexity in a world too complex to understand completely. It's a pragmatic necessity, not a natural kind. The boundaries we draw are ours, serving our

purposes, and we should change them when they stop serving those purposes. The world doesn't care about our boundaries. It just keeps changing, indifferent to whether our categories still fit.

9. Social Wholes and Moral Buck-Passing

In 2008, the global financial system nearly collapsed. Banks failed, retirement savings evaporated, millions lost their homes. When investigators tried to assign blame, they found something strange: nobody was responsible.

The traders were following incentives created by their compensation structure. The managers were maximizing shareholder value as their fiduciary duty required. The ratings agencies were using models approved by regulators. The regulators were enforcing laws passed by Congress. Congress was responding to constituents who wanted homeownership to be accessible. Every individual actor made locally reasonable decisions. The system produced catastrophe.

This is the problem of social wholes

doing things that none of their parts intended or chose. Groups act without minds. Institutions enforce rules nobody believes in. Systems generate outcomes nobody wanted. And when disasters occur, responsibility leaks between individuals and collectives in ways that let everyone off the hook simultaneously.

This chapter is about how parts hide inside wholes during both ordinary function and catastrophic failure, and why mereology is actually a moral problem in disguise.

Groups That Act Without Minds

A lynch mob is hundreds of people acting as one. No individual decides “we will hang this person.” The decision emerges from group dynamics: someone voices suspicion, others amplify it, emotion escalates, the crowd moves together toward violence. Each person feels swept along by the group’s

momentum. Nobody is “in charge.” The mob acts as a unified agent even though it has no unified decision-making process.

This should be impossible. Action requires intention. Intention requires a mind. Groups don’t have minds. Therefore groups shouldn’t be able to act. Yet groups clearly do act: armies invade, corporations merge, legislatures pass laws, protests topple governments.

The standard explanation is that groups are not really agents. They’re collections of individual agents whose actions happen to align. The alignment creates the appearance of group agency, but really it’s just coordinated individual agency. “The corporation decided” is shorthand for “various employees made decisions that, taken together, implemented a corporate strategy.”

This explanation works for well-

organized groups with clear decision procedures. The board votes, the CEO implements, the corporation acts. You can trace the action back to individual decisions that compose into the corporate action.

But it fails for emergent group actions where there's no clear decision procedure. The lynch mob has no vote. The financial crisis had no mastermind. The institution that systematically excludes minorities might have no explicitly racist individuals. The group behavior emerges from interaction patterns that nobody designed or controls.

This is mereological agency: the whole does things the parts didn't individually choose. The whole has causal powers beyond the sum of part-level powers. You can't explain the group action by listing individual actions because the

group action is a pattern of interaction, not a sum of isolated behaviors.

And if groups can act without minds, they can be responsible without any individual being responsible. The mob killed someone. Who is guilty? Everyone? No one? The people who threw the rope? The people who shouted loudest? The people who just watched? Responsibility doesn't neatly distribute to individuals because the action was collective.

Responsibility Leaking Upward and Downward

When something goes wrong in an organization, blame flows like water, finding the path of least resistance and pooling wherever it can hide.

Blame flows upward: "I was just following orders." The individual claims they were merely implementing

decisions made above. The manager made the decision, so the manager is responsible. The individual is just a part executing the whole's will. Responsibility stays with the whole.

Blame flows downward: "It was a rogue employee." The organization claims an individual acted outside policy. The organization has good policies and processes. One person violated them. The individual is responsible. The organization is a victim of part-level failure. Responsibility stays with the part.

Both moves are legitimate in some cases and self-serving in others. The trick is that mereological structure lets you make either move whenever it's convenient. When success happens, everyone claims credit. When failure happens, individuals blame the system and the system blames individuals. Responsibility disappears in the gap.

Consider a case: A bank forecloses on a home using forged documents. Who is responsible?

The document processor says: “I was processing thousands of documents monthly. I didn’t forge anything. I just signed where I was told to sign. How was I supposed to know they were fraudulent?”

The manager says: “I implemented the foreclosure targets set by senior leadership. I didn’t tell anyone to forge documents. I assumed the documents were legitimate. If someone cut corners to meet targets, that’s on them.”

The executives say: “We set reasonable performance goals. We had compliance policies. If employees violated policies to meet goals, that’s individual wrongdoing. We can’t monitor every document.”

The corporation says: “We’re committed

to following the law. This was a failure of individual employees who violated our policies. We've terminated the individuals responsible and strengthened our procedures."

Notice what happened: responsibility leaked between levels until it disappeared. Each level blames adjacent levels. Individuals blame the system for creating impossible targets. The system blames individuals for choosing illegal methods. Nobody is fully responsible because everyone can point to forces beyond their control.

The truth is probably that responsibility is distributed across levels. The individuals who forged documents are responsible for that choice. The managers who set unrealistic targets are responsible for creating conditions that incentivized fraud. The executives who ignored warning signs are responsible for

negligence. The corporate culture that prioritized speed over accuracy is responsible for systemic failure.

But distributed responsibility is psychologically and legally unsatisfying. We want someone to blame. We want punishment proportional to harm. Distributed responsibility means everyone is a little bit responsible, nobody is fully responsible, and the punishment never matches the harm.

“The Company Did It” as a Mereological Evasion

Corporations are legal persons with rights and responsibilities separate from their members. This is useful for commerce (the corporation can enter contracts that outlive any employee) and essential for limited liability (shareholders aren’t personally liable for corporate debts).

But legal personhood creates a mereological fiction: the corporation is both a whole composed of people and a distinct entity separate from those people. These two views are incompatible. If the corporation is just its members, then corporate action is member action and members should bear full responsibility. If the corporation is separate from its members, then it's not clear how it acts at all (it has no body or mind).

The law splits the difference. Corporations are distinct entities for some purposes (liability, taxation, contracts) and identical to their members for other purposes (criminal prosecution sometimes pierces the corporate veil). We switch between views depending on what's convenient.

This switching enables moral evasion. When the corporation does something

good, members claim credit (“we accomplished this”). When the corporation does something bad, members deflect blame (“the corporation did it, not me personally”).

“The company did it” is a way of hiding individual agency inside collective action. Individuals made decisions. The decisions composed into corporate action. The corporate action harmed people. But because the action is attributed to the corporation (a distinct legal entity), individuals escape responsibility. The corporation pays a fine. No person goes to jail. The harm is compensated but not punished.

This works because corporations can’t actually be punished. You can’t imprison a corporation. You can fine it, but fines are paid from corporate assets, which means they’re paid by shareholders, creditors, and future customers, most of

whom had nothing to do with the wrongdoing. The actual decision-makers often escape consequence entirely.

The mereological structure enables this. The whole (corporation) is treated as the agent and punished. The parts (individuals) escape because they were "just doing their jobs within a corporate role." The punishment falls on the wrong entity (a legal abstraction) while the responsible entities (actual humans who made decisions) face no consequences.

How Parts Hide Inside Wholes During Atrocities

The most disturbing cases are where ordinary people commit atrocities by fragmenting responsibility across roles and levels.

Hannah Arendt documented this in her study of Adolf Eichmann, the Nazi bureaucrat who organized transportation

to concentration camps. Eichmann didn't personally hate Jews. He wasn't ideologically committed to genocide. He was a logistics specialist optimizing train schedules. He saw himself as a competent administrator solving transportation problems.

Each step was locally reasonable. Receive an order: transport these people to these locations. Apply professional expertise: schedule trains, allocate resources, solve bottlenecks. Measure success: on-time performance, efficiency, cost reduction. None of these steps involved thinking about what happened at the destinations. That was someone else's role.

This is mereological evil: the action is monstrous at the whole level but unremarkable at the part level. Eichmann was a part in a genocidal whole, but he focused on his part (logistics) and

ignored the whole (genocide). The part-level work was morally neutral. The whole-level outcome was moral catastrophe.

This pattern repeats in every bureaucratic atrocity. The person approving loans didn't force anyone into homelessness. They processed applications according to policy. The person piloting the drone didn't murder civilians. They followed targeting procedures. The person designing the algorithm didn't create racial bias. They optimized for accuracy using available data.

Each part is defensible. The whole is indefensible. The parts compose into the whole, but the composition obscures responsibility because each part can honestly say "I didn't do anything wrong. I followed procedures. I did my job competently. The outcomes are not my fault."

This is why focusing on the system feels satisfying. The system is to blame. The individuals were trapped in roles that constrained their choices. Fix the system and you prevent future atrocities without having to blame individuals who were just following incentives.

But systems don't exist independently of the people implementing them. "The system" is an abstraction over patterns of behavior by actual humans. When you blame the system, you're displacing responsibility from people to a pattern. The pattern can't be punished. The people can, but we just said they're not responsible because the system made them do it.

The circle is vicious. Individuals create systems. Systems constrain individuals. Individuals claim the system made them do it. Systems can't be responsible because they're not agents. Everyone

escapes responsibility by pointing at something else.

When Parts Can't Be Excused by the Whole

Yet some parts are responsible even when embedded in wholes. The question is when.

One answer: when they have discretion. If your role is pure execution with no judgment, you're not responsible for outcomes. But if you exercise judgment, make decisions, or have discretion about implementation, you bear some responsibility for how you use that discretion.

Eichmann had discretion. He could have requested reassignment. He could have worked less efficiently. He could have sabotaged the operation. He chose not to. The choice was constrained by career incentives, social pressure, and political

danger, but it was still a choice. He bears responsibility for that choice even though he was part of a larger genocidal apparatus.

Another answer: when you benefit. If you profit from wrongdoing, you're responsible even if you didn't directly cause it. The shareholders who profited from corporate fraud are responsible even if they didn't forge documents. They benefited from the fraud and had power to prevent it (through governance rights), so they bear responsibility for not exercising that power.

A third answer: when you enable. If your part is essential to the whole's wrongdoing, you're responsible even if you're not making decisions. The person who maintains the torture equipment is responsible even if they never tortured anyone. Without their maintenance, torture wouldn't be possible. They enable

the whole's action and are therefore complicit.

These principles are contestable and often in tension. Eichmann had discretion but faced serious constraints. Shareholders benefit but lack detailed knowledge. Equipment maintainers enable but may have no alternatives (quit and someone else takes the job). There's no clean formula for distributing responsibility between parts and wholes.

What's clear is that you can't automatically excuse parts by pointing to wholes. "I was following orders" is sometimes a valid defense (when orders leave no discretion) and sometimes an evasion (when discretion existed but wasn't used to resist). "The system made me do it" is sometimes true (when alternatives don't exist) and sometimes false (when alternatives exist but are costly).

The mereological question “am I responsible for what the whole does?” doesn’t have a general answer. It depends on your causal contribution, your discretion, your knowledge, your power to resist, and your moral imagination. These vary by role, context, and individual capacity.

Collective Responsibility Without Collective Agents

Can groups be responsible when they’re not really agents? Can unorganized collections bear responsibility for outcomes they collectively produced?

Climate change is the paradigm case. No individual is responsible. No nation is solely responsible. No company is responsible. Yet humans collectively are responsible. We emit greenhouse gases. The emissions accumulate. The planet warms. Disasters follow. The chain of causation is clear. The responsibility is

real. But it's distributed across billions of people making ordinary choices (driving, flying, eating meat, using electricity) that individually seem harmless.

How can responsibility be collective without being assignable to individuals? Standard moral theory struggles here. Responsibility requires agency. Agency requires intention. No individual intends climate change. The collection of individuals isn't an agent. So who is responsible?

One answer: everyone who contributes is partially responsible, weighted by contribution and capacity to reduce it. The billionaire flying private jets is more responsible than the subsistence farmer. The person who could easily reduce emissions but doesn't is more responsible than the person with no alternatives. Responsibility is distributed proportionally.

This answer is logical but practically useless. It makes everyone a little bit responsible, nobody very responsible, and provides no clear target for blame or punishment. It also faces the “drop in the ocean” problem: my individual contribution is negligible, so how can I be responsible for a global catastrophe?

Another answer: responsibility attaches to agents capable of addressing the problem, regardless of causation. Rich nations are responsible not because they caused all the emissions (though they caused disproportionate shares), but because they have capacity to fix it. With great power comes responsibility, even for problems you didn’t create.

This answer justifies international climate agreements where rich nations commit to larger emission reductions and pay for adaptation in poor nations. But it decouples responsibility from causation,

which feels unfair. Why should you be responsible for fixing problems you didn't cause?

A third answer: reject the individual-level framing entirely. Climate change is a system-level problem requiring system-level solutions. Asking who is individually responsible is the wrong question. The right question is: what institutional changes would reduce emissions? The answer might be carbon taxes, regulations, international agreements, technological investment. These are collective decisions implemented through institutions, not individual moral choices.

But institutions are composed of individuals. Passing the buck from individuals to institutions just moves the question: which individuals in which institutions bear responsibility for implementing solutions? We're back to

mereological buck-passing.

Institutional Capture and Distributed Corruption

Sometimes institutions are corrupted not by corrupt individuals but by selection effects and incentive alignment that produce institutional outcomes nobody would individually endorse.

Regulatory capture is the standard example. An agency meant to regulate an industry becomes controlled by that industry. This doesn't require conspiracy or bribery. It happens through normal processes: industry hires former regulators (they have expertise), regulators defer to industry (they have information), the revolving door aligns incentives (today's regulator is tomorrow's consultant), and gradually the agency serves industry interests rather than public interests.

No individual is corrupt. Each person is making locally reasonable career choices. Accepting a high-paying industry job after government service is legal and rational. Consulting with industry experts is good policy. Deferring to technical expertise is prudent. But the aggregate effect is that the institution stops functioning as intended.

This is mereological corruption: the whole is corrupted even though the parts are clean. The corruption is structural, not personal. You can replace all the individuals and the corruption persists because the structure creates the same incentives for the new individuals.

Fixing this requires structural change: stronger revolving-door restrictions, independent funding, transparency requirements. But structural change is politically difficult because the corrupted institution fights it (using its captured

power) and individual members resist (they're not corrupt, why are we being punished?).

The mereological problem is that structural problems require structural solutions, but structures are composed of individuals, and individuals resist changes that harm their interests even when those changes would improve the whole's function.

The Upshot

Social wholes do things their parts don't intend. Responsibility leaks between levels. Parts hide inside wholes during both ordinary function and catastrophic failure. And mereology provides the vocabulary for these dynamics without providing solutions.

The lesson is not that individuals are never responsible. It's that individual responsibility exists alongside collective

responsibility, and neither fully explains outcomes in complex social systems. Atrocities require individual choices and systemic failures. Reform requires individual courage and structural change. You can't reduce everything to individuals or everything to systems.

The error is trying to pick one level as "the real" level of moral responsibility. Both levels are real. Both matter. Individuals make choices within constraints created by systems. Systems are sustained by patterns of individual behavior. Change requires intervention at both levels simultaneously.

When institutions fail catastrophically, the right response is not to decide whether individuals or systems are to blame. It's to trace how individual decisions composed into system-level failures, identify where responsibility lies at each level, and intervene at the levels

where intervention is most effective.

Sometimes that means prosecuting individuals who exercised discretion wrongly. Sometimes it means restructuring institutions to change incentives. Often it means both. The mereology doesn't tell you which. It just reveals the structure of the problem: parts and wholes, each with causal powers, each capable of producing harms that can't be fully explained at either level alone.

Social wholes are where mereology stops being abstract philosophy and becomes a practical tool for understanding how humans organize into groups that do things individuals wouldn't do alone. The tool doesn't make the moral questions easier. It makes them clearer. That's the best philosophy can do.

10. Who Is Responsible When Parts Misbehave?

A self-driving car kills a pedestrian. Who is responsible?
The car's sensors detected the pedestrian but classified them as a non-threat (maybe a plastic bag blowing across the road). The classification algorithm was trained on data that underrepresented this pedestrian type. The training data was collected by a contractor following specifications from the car company. The specifications were written by engineers optimizing for cost and speed. The executives approved the specifications to meet a product launch deadline. The regulators approved the car based on testing that didn't catch this failure mode.

Every person in this chain made locally reasonable decisions within their role. The sensor worked correctly. The

algorithm performed as designed. The contractor met specifications. The engineers balanced competing constraints. The executives managed business pressures. The regulators applied standard procedures. Yet someone is dead.

Who is responsible? This is not an abstract philosophical question. It's a practical legal question that courts will have to answer, and the answer will determine whether autonomous vehicles are viable as a technology. But it's also a moral question that cuts to the core of what responsibility means in systems where outcomes emerge from the composition of many parts.

This chapter is about the hardest problem in mereological ethics: when parts compose into wholes that cause harm, how do you distribute responsibility?

Individual Guilt vs Systemic Blame

Start with the easy case: a drunk driver kills a pedestrian. The driver is responsible. They chose to drink. They chose to drive. The choice directly caused the death. We assign full responsibility to the individual, prosecute them, and feel justice has been served.

Now the hard case: a sober driver, following all traffic laws, driving a car with defective brakes that the manufacturer knew about but didn't recall, hits a pedestrian who stepped into traffic without looking. Who is responsible?

The driver didn't choose defective brakes. The manufacturer chose profit over safety. The pedestrian chose to step into traffic. The city chose not to install a crosswalk at a dangerous intersection. Responsibility is distributed across multiple agents making independent decisions that composed into a fatal

outcome.

The law handles this through negligence: each party is responsible to the degree they failed to exercise reasonable care. The manufacturer is liable for defective brakes. The pedestrian might bear partial responsibility for jaywalking. The city might be liable for unsafe road design. The driver might be blameless if they couldn't have avoided the accident even with functioning brakes.

This works okay when you have clear causal chains and identifiable failures of care. It breaks down when you have systemic failures where nobody is clearly negligent but the outcome is still catastrophic.

Consider the 2008 financial crisis again. Who was negligent? The traders were compensated for taking risks, so they took risks. Were they negligent for responding to incentives? The banks

were maximizing shareholder value as fiduciary duty requires. Were they negligent for doing what they're legally obligated to do? The regulators were applying rules written by Congress. Were they negligent for following the law?

You can argue everyone was negligent (they should have seen the systemic risk), or nobody was negligent (they were following local incentives and rules). Either way, negligence doesn't capture what went wrong. What went wrong was that individually rational decisions composed into collective irrationality. The system produced outcomes nobody wanted and nobody could prevent through individual action.

When systemic failures occur, blaming individuals feels like missing the point. The individuals were interchangeable. Replace them with equally competent people and you get the same outcome.

The problem is the system, not the people.

But systems don't exist independently of people. Blaming "the system" is blaming a pattern of behavior by actual humans. If the humans aren't responsible for their behavior, who is?

This is the tension: individual responsibility makes sense when individuals have meaningful choice. Systemic responsibility makes sense when outcomes emerge from structural constraints that overwhelm individual choice. Most real cases involve both, and there's no formula for weighing them.

Why Blaming "the System" Feels Satisfying

Systemic blame is psychologically comfortable because it lets everyone off the hook simultaneously. Nobody is a bad person. Nobody made evil choices.

The system is broken. Fix the system and the problem goes away. Everyone can feel they did their best within a bad situation.

This is particularly appealing to people caught in the system. Of course they blame the system. They don't feel responsible for outcomes they didn't intend and couldn't prevent. They were doing their jobs, following procedures, meeting expectations. The outcomes are not their fault.

But systemic blame also appeals to observers. Blaming individuals requires identifying who specifically did wrong and proving they had realistic alternatives. This is hard work. It requires detailed understanding of each person's situation, options, and constraints. It opens you to accusations of scapegoating or failing to understand the full context.

Blaming the system is easier. You don't need to identify individual culprits. You don't need to prove anyone had malicious intent. You just point at bad outcomes and say "the system produced this, the system is broken, we need systemic reform." This sounds sophisticated and avoids the appearance of witch-hunting.

The problem is that systemic blame without individual accountability creates moral hazard. If nobody is personally responsible for systemic failures, people have no incentive to resist systemic pressures. Why sacrifice your career to resist a policy you think is harmful if you can later say "I was just working within the system"? The lack of accountability enables the very behavior that sustains broken systems.

Effective accountability requires both levels. Yes, systems create constraints

and incentives. But individuals choose how to respond to those constraints. Some people resist systemic pressures at personal cost. They're heroes or whistleblowers depending on whether they succeed. Others go along with systems they know are harmful. They're complicit.

The distinction matters. If everyone is equally constrained by the system, nobody can resist, and "resistance was impossible" excuses all behavior. But actually some people do resist. They face costs, but resistance is possible. The people who don't resist bear responsibility for that choice, even when the choice is hard.

When Parts Can't Be Excused by the Whole

There are cases where appealing to the whole doesn't excuse the part. When are those cases?

First, when you have discretion and use it wrongly. A police officer shooting an unarmed person can't excuse it by saying "I was trained to perceive threats." The training might have contributed, but the officer had discretion in that moment and used it to kill someone who wasn't actually threatening them. The systemic problems with police training don't eliminate individual responsibility for individual decisions.

Second, when you could have exited or resisted. A mid-level manager implementing a fraudulent scheme can't fully excuse themselves by saying "I would have been fired if I refused." Getting fired is bad. Participating in fraud is worse. You chose fraud over unemployment. That's a choice you're responsible for.

The counterargument is that this sets an unreasonably high bar for resistance.

Most people can't afford to lose their jobs. Expecting them to sacrifice their careers for moral principles is demanding martyrdom. Systemic change shouldn't require individual heroism.

This is true but incomplete. Yes, systemic pressure is real. Yes, most people can't afford to resist. But some people are positioned to resist at lower cost, and they bear greater responsibility for not resisting. The executive who could blow the whistle and find another executive job is more responsible than the low-level employee who would become unemployable if they resisted.

Responsibility scales with power and alternatives. The more power you have, the more responsible you are for how you use it. The more alternatives you have, the less you can excuse bad choices by claiming constraint.

Third, when you benefit. If you profit

from wrongdoing, you're responsible even if you didn't cause it. The shareholders of a company committing fraud are responsible even if they didn't know about the fraud. They benefited, they had power (through voting and governance), and they failed to exercise oversight. Benefit plus power equals responsibility.

Fourth, when you enable. If your contribution is essential to the harm, you're responsible even if you're not making decisions. The accountant who cooks the books is responsible even if executives ordered it. The engineer who designs the torture device is responsible even if they never operated it. Without your specialized contribution, the harm wouldn't be possible. You enabled it. You're complicit.

These principles overlap and conflict. A low-level employee might enable harm

(necessary contribution) but have no alternatives (would be fired and unemployable) and not benefit (just collecting a salary). How responsible are they? There's no formula. You have to weigh the factors case by case.

What's clear is that "I was just a part in a larger system" doesn't automatically excuse you. Parts bear responsibility based on their causal contribution, discretion, alternatives, power, and benefits. These vary enormously across parts. Responsibility varies accordingly.

The Moral Relevance of Scale

Responsibility intuitions change with scale in strange ways. Kill one person and you're a murderer. Structure a financial system that leads to thousands of deaths through suicide, stress-related illness, and lack of medical care due to poverty, and you're just an economist who got some predictions wrong.

This asymmetry is weird. The economist's actions caused more deaths than the murderer's. The deaths are more indirect, but they're causally downstream of decisions the economist influenced. If causation matters for responsibility, the economist should be more responsible, not less.

But it feels wrong to equate them. The murderer intended to kill. The economist didn't intend deaths, just economic efficiency. Intent matters morally. The murderer is more culpable even if the economist caused more harm.

Or consider: a factory knowingly pollutes a river, causing cancer in downstream communities. Hundreds die over decades. The factory executives are rarely prosecuted, and when they are, they face fines and maybe short prison sentences. Compare this to a serial killer who murders five people. Life in prison,

universal condemnation, no moral ambiguity.

Again, the scale of harm is greater from the factory. But the moral intuition is that the serial killer is worse. Why?

Part of it is intent. The serial killer intended each death. The executives intended profit, and deaths were a side effect. We treat intended harm as morally worse than caused harm, even when caused harm is quantitatively greater.

Part of it is directness. The serial killer directly killed people. The executives set in motion causal chains that eventually led to deaths. We treat direct causation as morally weightier than indirect causation.

Part of it is distributedness. The serial killer is a single agent. The factory deaths involve multiple agents: executives who approved pollution, engineers who designed inadequate

controls, regulators who failed to enforce standards, politicians who wrote weak laws. Responsibility is distributed, so blame is diluted.

These intuitions might be tracking something important: intent and directness are signals of character. The serial killer has a defective moral character in a way the negligent executive might not. The executive might be greedy and callous but not sadistic. The distinction matters if you care about moral character rather than just outcomes.

But it creates a perverse incentive: structure your harmful activities to be indirect, distributed, and plausibly unintended, and you can cause massive harm while facing minimal moral and legal consequences. This is exactly what corporations do. Harm is scaled, distributed across divisions and

subcontractors, and structured to be plausibly unintended. It's a mereological strategy for evading responsibility.

Responsibility in Distributed Systems

Now take the hardest case: algorithmic decision-making. An algorithm denies someone a loan. The denial causes financial hardship, maybe even death (stress-related illness, inability to afford medical care). Who is responsible?

The algorithm is not responsible (it's not an agent). The programmers might not be responsible (they wrote general code, not specific to this case). The company deploying the algorithm might not be responsible (they're using industry-standard tools). The person denied the loan might be partially responsible (they have bad credit for reasons that might be their fault). But someone suffered serious harm, and it feels like someone should be responsible.

Distributed responsibility across the system: the programmers are partially responsible for writing biased code, the company is partially responsible for deploying it without adequate testing, the regulators are partially responsible for not requiring audits, society is partially responsible for creating conditions where credit score predicts race, the applicant is partially responsible for past financial decisions.

This distribution is logical but unsatisfying. It makes everyone a little responsible and nobody very responsible. The harm is real and concentrated (one person suffered). The responsibility is diffuse and diluted (many people contributed small amounts). The person harmed can't seek redress because there's no clear target for blame.

This is the mereological structure of algorithmic harm: causation is distributed

across many agents making independent decisions. The harm emerges from composition of those decisions. No individual decision is clearly wrong (each step is locally justifiable). The composition is catastrophic.

Traditional responsibility frameworks assume individual agents causing direct harm through intentional action. Algorithmic systems violate all these assumptions: the agent is a system, the causation is indirect, the harm is a statistical outcome across many decisions, nobody intended it. Our moral concepts don't fit.

The pragmatic solution is to assign responsibility to whoever has power to change the system. If the company deploying the algorithm can audit it, modify it, or stop using it, the company is responsible. Not because they caused the specific harm, but because they can

prevent future harms and they're profiting from the system.

This decouples responsibility from causation. You're responsible not because you caused harm, but because you have power to prevent harm and you're not using that power. This is a different kind of responsibility: not guilt for past actions, but obligation for future actions.

It's also pragmatically necessary. In complex systems, tracing responsibility to specific causal contributions is impossible. There are too many contributors, too many feedback loops, too many emergent effects. The only workable approach is to identify who has power to fix problems and assign them obligation to use that power.

Collective Punishment and Moral Pollution

When you can't identify individual culprits, collective punishment is tempting. The company committed fraud, so fine the company. The institution is racist, so sanction the institution. The nation committed atrocities, so impose reparations.

Collective punishment makes sense when the collective is genuinely responsible. If a company systematically defrauds customers through corporate policy, the company (as an entity) is responsible. Punishing it sends the right signal.

But collective punishment creates problems. The punishment falls on current members who might not have participated in the wrongdoing. Fining a company punishes current shareholders (who might be pension funds investing on behalf of retirees). Sanctioning an institution punishes current members (who might have joined recently and

opposed the bad behavior). Reparations punish current taxpayers (who might not have benefited from historical injustices).

Is this fair? The people being punished didn't do the wrong thing. They just happen to be members of an entity that did wrong things. We're punishing parts for whole-level wrongdoing they didn't contribute to.

One response: membership creates responsibility. If you join an institution with a history of wrongdoing, you inherit some responsibility for addressing that wrongdoing. You benefit from the institution's resources and reputation (built partly through past wrongdoing), so you bear some burden for making amends.

This is the logic of reparations. Current citizens didn't enslave anyone, but they inherit a nation built partly on slavery. They benefit from accumulated wealth

that includes wealth extracted through slavery. They bear some obligation to address continuing harms from that history.

Another response: collective punishment is not about fairness to current members. It's about incentive effects. If institutions can commit wrongdoing and escape punishment by replacing members, they'll have no incentive to prevent wrongdoing. Collective punishment creates incentives for institutions to police themselves and for individuals to avoid joining institutions with bad track records.

Both responses justify some collective punishment, but they don't resolve the unfairness to individuals who bear punishment for wrongs they didn't commit. That unfairness is the cost of maintaining institutional accountability. You can have institutional accountability

with some unfairness to individuals, or you can have perfect individual fairness with no institutional accountability. You can't have both.

Legal Fictions and Moral Realities

Corporations are legal persons. This fiction is essential for modern commerce. It lets corporations enter contracts, own property, and sue/be sued. It creates continuity despite changing membership. It enables limited liability, which encourages investment.

But the fiction creates moral confusion. If corporations are persons, they should be morally responsible like persons. But they're not really persons. They're organizations. Organizations don't have intentions or feelings or moral agency in the way persons do.

The law navigates this by treating corporations as persons for some

purposes (contract, property) and as organizations for others (criminal prosecution sometimes pierces the corporate veil to prosecute individuals). We switch between treating them as wholes (liable as entities) and as collections of parts (individual managers can be prosecuted).

This switching is pragmatic but incoherent. Either corporations are moral agents responsible for their actions, or they're just collections of individuals and only individuals can be responsible. The law wants both: corporate liability when it's convenient (deep pockets for damages) and individual liability when corporate liability seems insufficient (prosecution of executives for fraud).

The incoherence reflects a genuine mereological problem: wholes really do things their parts don't individually do. But wholes don't have the psychological

properties (intentions, beliefs, feelings) we usually associate with moral agency. So we have entities that act and cause harm but don't fit our paradigms of moral responsibility.

The practical solution is to treat corporate responsibility as different from individual responsibility. Corporations can be liable (pay damages) without being guilty (morally culpable). Liability addresses harm and creates incentives. Guilt is reserved for agents with intentions, which means individuals.

This works legally but creates a gap morally. Corporate harms often have no guilty individuals (everyone acted reasonably within their roles). We can assign liability but not guilt. Justice requires both punishment (for the guilty) and compensation (for the harmed). We get compensation without punishment, which feels incomplete.

The Upshot

Responsibility in mereologically complex situations is irreducibly messy. You have wholes that act through parts, parts that hide in wholes, distributed causation, emergent outcomes, and no clear answer to who is really responsible when things go wrong.

The traditional approaches fail. Pure individualism (only individuals are responsible) can't explain systemic failures where individuals are interchangeable. Pure collectivism (only systems are responsible) eliminates accountability and enables moral hazard. Both levels matter, and the right distribution of responsibility varies by case.

The principles that help:

- Responsibility scales with discretion, power, alternatives, and benefits

- Intent and directness matter morally, not just causal contribution
- Collective responsibility is real but distributes unfairly to individuals
- Power to prevent future harm creates obligation independent of causation
- Different contexts require different levels of analysis (individual for character judgments, systemic for structural reform)

These principles don't give you a formula. They give you factors to weigh. The weighing requires judgment, and reasonable people will disagree. That's inevitable when dealing with complex systems where causation and responsibility don't align neatly.

What's not inevitable is pretending the problem doesn't exist or defaulting to easy answers (blame individuals, or blame the system, but not both). The

problem is real. The complexity is inherent to mereological structure. The best you can do is be honest about the factors at play and make principled judgments about how to weight them.

This is not satisfying. It doesn't give closure. It means living with moral uncertainty about who is really responsible when parts compose into wholes that cause harm. But moral certainty in these cases is a sign you're ignoring complexity, not that you've solved the problem. The complexity is real. The uncertainty is appropriate. The task is to act justly despite the uncertainty, not to pretend the uncertainty away.

11. States, Institutions, and Abstract Beasts

The United States government is the largest employer on Earth. It commands nuclear weapons, regulates commerce, imprisons millions, and extracts trillions in taxes. It does all this despite being composed entirely of temporary parts that are constantly replaced and that mostly can't stand each other.

Presidents serve four or eight years. Members of Congress turn over. Federal employees retire. The military cycles personnel. The only constant is the structure: the Constitution, the laws, the organizational chart, the procedures. The structure persists while the people flow through it like water through a pipe. The pipe does things. The water just moves.

This is the fundamental strangeness of institutions: they're real entities with

causal powers, but they exist almost entirely as abstract structure. Remove all the people and the institution disappears. But the people are replaceable. The institution persists through complete personnel replacement. What kind of object is real, powerful, and independent of its material substrate?

This chapter is about the strangest wholes of all: the ones that exist primarily as patterns in behavior sustained by collective belief, and that might be more powerful than any physical object you'll encounter.

Governments as Part-Machines

Think of government as a machine for converting individual actions into collective decisions. The machine has components: legislatures that convert preferences into laws, executives that convert laws into policies, bureaucracies that convert policies into actions, courts

that convert disputes into judgments. The components are made of people, but the people are interchangeable. What matters is the role, not the person filling it.

This makes governments ontologically weird. A car is a machine made of metal parts. Remove the engine and the car stops being a car. Replace the engine and it's still the same car (usually). But governments are machines where the parts are people, and people have agency. They can refuse to play their roles. They can reinterpret their roles. They can coordinate to change the machine's structure. The parts can rebel against the whole in ways engine pistons cannot.

Yet governments mostly work. People mostly play their assigned roles. Legislators legislate. Bureaucrats implement. Judges judge. They do this despite having individual interests that conflict with their institutional roles.

Why?

Part of the answer is incentives. You're paid to perform the role. You face sanctions for failing to perform it. You gain status and career advancement from performing it well. The incentives align individual behavior with institutional function most of the time.

But incentives can't be the whole answer because they don't explain why people care about institutional legitimacy. A purely incentive-driven bureaucrat would take bribes whenever the benefit exceeds the risk of punishment. Many don't, even when they could get away with it. They care about doing the job properly, not just avoiding punishment.

This suggests something beyond incentives: internalization of institutional norms. You don't just play the role because you're paid to. You identify with the role. You adopt its standards as your

own. You feel pride in doing it well and shame in doing it poorly. The institution has colonized your psychology.

This colonization is how institutions sustain themselves. They recruit people, train them in the role, reward performance, and gradually transform them into parts that serve institutional function. The people remain agents with their own interests, but they've partially aligned their interests with the institution's. The alignment is never perfect, but it's enough for the institution to function.

The result is a machine made of partially aligned agents rather than purely mechanical parts. The machine is more flexible than a physical machine (it can adapt to novel situations through human judgment) but more fragile (the human parts can defect or malfunction in creative ways). Managing this fragility is

what institutional design is about.

Laws as Binding Agents

If institutions are machines, laws are the binding agents that hold the parts together. Not physically bound (people can walk away from their jobs), but normatively bound. The law says what ought to happen, and people mostly comply, creating stable patterns of behavior.

But what makes law binding? Not physical force. The government has force (police, military), but it can't physically compel all behavior. There aren't enough police to ensure everyone pays taxes or obeys speed limits. Compliance has to be mostly voluntary.

One answer: laws bind because they're backed by sanctions. You obey because you don't want to be punished. This explains some compliance but not all.

Many people obey laws even when the chance of punishment is negligible. You don't steal from unlocked cars even when nobody is watching. You don't jaywalk even when no cop is around. The behavior isn't fear of sanctions.

Another answer: laws bind because they're legitimate. You obey because you believe the law has authority, which means it has the right to command your behavior. Where does this authority come from? Democratic consent, constitutional procedure, moral correctness, or just sheer social acceptance. The source varies, but the effect is the same: you feel obligated to obey.

This obligation is the binding force. It's not physical (you could disobey) or purely incentive-based (sometimes obedience is costly). It's normative: you ought to obey, and you mostly want to do

what you ought to do. The “ought” is enforced partly by external sanctions and partly by internal commitment to playing your part in a legitimate order.

This makes law a strange kind of glue. It binds people into institutional structures through a combination of force, incentives, and internalized obligation. The proportion varies: totalitarian states rely more on force, liberal democracies rely more on legitimacy, most systems use a mix. But all functional legal systems create binding through more than just physical coercion.

The binding is also self-reinforcing. Laws work because people obey them. People obey them because they work (they coordinate behavior, resolve disputes, enable cooperation). The working legitimizes the laws, which increases obedience, which makes them work better. It’s a virtuous cycle when it

functions, a vicious cycle of delegitimization when it breaks down.

Revolutions as Re-Partitioning Events

Sometimes the binding breaks. The laws lose legitimacy. People stop obeying. The institutional structure collapses. This is revolution: the rapid dissolution of one institutional order and its replacement with another.

Revolutions are mereological catastrophes. The whole disintegrates into parts that refuse to play their roles. The police won't enforce. The military won't suppress protests. The bureaucrats won't implement orders. The tax collectors won't collect. Each individual decides the institution no longer has authority, and once enough individuals defect, the institution ceases to exist.

What's strange is how quickly this can happen. An institution that seemed stable

for decades can collapse in weeks. The Soviet Union appeared permanent. Then it dissolved in a matter of years, with surprisingly little violence. The parts stopped believing in the whole, and the whole evaporated.

This reveals that institutions are sustained by collective belief in a way physical objects are not. A building doesn't stop being a building if people stop believing in it. But a government does stop being a government if enough people stop recognizing its authority. The institutional structure is real and causally powerful, but it depends on ongoing collective recognition in ways that physical structures don't.

Revolutions don't just dissolve old institutions, they create new ones. The revolutionary moment is when old institutional boundaries dissolve and new ones can be drawn. Who counts as a

citizen? What counts as legitimate authority? How should power be organized? These questions are usually settled, but revolution unsettles them.

The re-partitioning is often violent because different groups want different boundaries. Should the new state include this ethnic group or exclude them? Should power be centralized or distributed? Should property rights be preserved or abolished? These aren't just policy questions. They're questions about what kind of institutional structure will exist and who will be inside or outside it.

The violence comes from the fact that institutional boundaries determine who has power, rights, and resources. Changing the boundaries redistributes all three. The people who benefit from current boundaries defend them. The people who would benefit from different boundaries attack them. The conflict is

over the fundamental structure of the social whole, not just over policies within that structure.

Once new boundaries are established and enough people accept them, the new institution solidifies. But the solidification takes time. The new state has to prove it can govern (provide services, maintain order, defend borders). It has to cultivate legitimacy (convince people it has authority). It has to bind the parts into a functional whole (create incentives and norms that align behavior with institutional function). If it succeeds, the revolutionary re-partitioning becomes the new stable structure. If it fails, you get prolonged instability or another revolution.

Why Reform Fails When Parts Are Misidentified

Most institutional change is not revolutionary. It's reform: modifying the

structure while preserving basic continuity. This requires correctly identifying which parts need changing and which need preserving.

The identification is harder than it looks. Institutions are complex systems with many interacting parts. Change one part and you create ripple effects through the whole system. Sometimes the effects are intended. Often they're not.

Example: welfare reform in the 1990s aimed to reduce dependency by adding work requirements. The theory was that requiring work would give people incentive to become self-sufficient. The reform changed the incentive structure (one part of the welfare system) while leaving other parts mostly intact (benefit levels, eligibility, administration).

The reform had the intended effect (welfare rolls declined) and unintended effects (many people fell into extreme

poverty because they couldn't find steady work that paid enough to replace benefits). The reformers misidentified which parts were causing dependency. They thought it was incentives. Actually it was a combination of low wages, lack of childcare, transportation barriers, health problems, and discrimination. Changing incentives without addressing these other factors just pushed people off welfare without enabling self-sufficiency.

This is a classic mereological failure: treating a system-level problem (poverty) as if it could be solved by adjusting one part (welfare incentives) while ignoring other parts that constrained what welfare recipients could actually do. The parts weren't misidentified individually (welfare incentives are real), but their causal role in the whole was misunderstood.

Reform fails when you misidentify which

parts are producing the outcomes you want to change. Sometimes you're changing the wrong part (welfare incentives instead of wage structures). Sometimes you're treating a whole-level problem as a part-level problem (systemic racism as individual prejudice). Sometimes you're ignoring how parts interact (changing one part breaks integration with other parts).

Successful reform requires understanding the system as a whole, identifying where intervention will be most effective, and anticipating how changes propagate through the system. This is rare because most reformers focus on their area of expertise, which is usually one part of the system. They optimize locally without understanding global effects.

Bureaucracies as Self-Perpetuating Composites

Bureaucracies are institutions optimized

for persistence, not performance. They survive by making themselves indispensable and by distributing power so widely that no individual can dismantle them.

The classic bureaucratic structure is hierarchical: levels of authority from top to bottom, with each level controlling the level below. But the real power structure is more complex. Each level has information the levels above and below don't have. Each level can slow-walk, reinterpret, or selectively implement directives. The nominal hierarchy is overlaid with a network of informal power based on expertise, relationships, and control of resources.

This makes bureaucracies nearly impossible to reform from the top down. The executive issues a directive. Middle management "implements" it in ways that preserve their interests. Front-line

workers adapt in ways that minimize disruption to their routines. The directive gets filtered through layers of interpretation until it bears little resemblance to what was intended. The executive claims success. The bureaucracy remains unchanged.

The persistence mechanism is distributed self-interest. Each part of the bureaucracy has an interest in preserving itself. Job security, budgets, status, and power all depend on the part continuing to exist. So each part resists changes that would eliminate or reduce it. No part can resist successfully alone, but the distributed resistance across many parts is very effective at blocking reform.

This is why bureaucracies are so stable. They're composed of parts that individually are replaceable (any bureaucrat can be fired) but collectively are entrenched (the whole bureaucracy

can’t be dismantled without political will that exceeds what reformers can usually muster). The whole persists through the self-interested behavior of parts that each have weak power individually but strong power collectively.

The same mechanism that makes bureaucracies persistent makes them sclerotic. They’re optimized for survival, not adaptation. They resist changes that would make them more effective if those changes threaten any part’s interests. The result is institutions that persist long past their usefulness, consuming resources without producing proportional value.

Breaking this requires either external shock (crisis that delegitimizes the bureaucracy and enables radical reform) or patient internal reform that builds coalitions across parts and creates incentives for change. Both are difficult. Most bureaucracies outlive their critics.

Corporations as Legal Fictions With Real Power

Corporations are even stranger than governments. They're legally persons with rights, but everyone knows they're not really persons. They're organizational structures given personhood by legal fiat. The fiction enables limited liability (shareholders aren't personally liable for corporate debts) and perpetual existence (the corporation persists through ownership changes).

But the fiction creates real power. Corporations can accumulate capital across generations. They can enter contracts that bind future owners. They can own property, sue, and be sued. They're more stable than individuals (they don't die) and more powerful than most individuals (they have more resources). The legal fiction has become

a dominant actor in modern economies and politics.

This is mereology at its strangest: a purely legal entity with no physical existence beyond its parts (shareholders, employees, assets) but with causal powers that exceed any individual part. The corporation does things. It builds products, provides services, influences politics, shapes culture. None of this is done by "the corporation" as a physical entity. It's done by people acting in corporate roles, coordinated by corporate structure, pursuing corporate interests.

But corporate interests aren't just the sum of individual interests. Shareholders want profit. Employees want wages and job security. Customers want products. The corporation balances these interests according to legal rules and internal governance. The balance produces corporate behavior that often doesn't

match what any stakeholder would individually choose.

This is clearest in cases of corporate malfeasance. The corporation commits fraud or causes environmental damage or exploits workers. Who chose this? Often nobody. The CEO is maximizing shareholder value as fiduciary duty requires. The managers are meeting metrics set by the CEO. The employees are following orders. The shareholders are passive investors who don't know details of corporate operations. The behavior emerges from the interaction of roles and incentives without anyone specifically choosing it.

This emergence of corporate behavior from part-level actions without whole-level intention creates the accountability gap discussed earlier. The corporation did it, but the corporation is a legal fiction. The people did it, but they were

playing roles in a structure that constrained their choices. Responsibility leaks between the legal fiction and the actual humans until it disappears.

The solution is to treat corporate responsibility as sui generis: different from individual responsibility but still real. Corporations are responsible in the sense that they can be held liable (forced to pay damages, change behavior, or dissolve). Individuals within corporations are responsible when they exercise discretion wrongly or enable wrongdoing. Both levels of responsibility are necessary because both levels of agency are real.

Constitutional Versus Ordinary Politics

Constitutions are meta-level institutions: rules about rules, structure about structure. They define how ordinary political decisions get made. They're

supposed to be stable (changed only through difficult amendment processes) while ordinary politics is fluid (laws change through normal legislative processes).

This creates a two-level system: constitutional politics determines the game, ordinary politics plays within the game. The separation is meant to provide stability (the game doesn't change every time someone wins) while allowing adaptation (rules within the game can change as needed).

But the separation is always contested. Groups dissatisfied with the game want to change it. Groups benefiting from the game want to preserve it. The fight is about whether an issue is constitutional (requiring fundamental change to how decisions are made) or ordinary (solvable within existing procedures).

Example: voting rights. Is the question

“should this category of people vote?” an ordinary political question (decide by normal legislative process) or a constitutional question (requiring amendment or judicial interpretation of fundamental rights)? The answer determines how hard it is to change. Ordinary questions require a legislative majority. Constitutional questions require supermajorities, judicial review, or formal amendment.

The level at which a question is fought is as important as the answer. Treating something as constitutional elevates it beyond normal politics. It makes the rule harder to change and (if you win) more stable. Treating something as ordinary politics leaves it vulnerable to shifts in legislative majorities.

This is a mereological question about which level of the system is appropriate for which decision. Some things should

be stable (fundamental rights, basic procedures, allocation of power between institutions). Some things should be fluid (budgets, regulations, specific policies). Getting the level wrong produces either excessive rigidity (too many things constitutionalized) or excessive instability (too few things protected from ordinary politics).

Most constitutional systems don't clearly specify which things belong at which level. They give general principles and let courts interpret them. This makes the boundary between constitutional and ordinary politics itself a political question. The fight is often about the boundary, not about specific policies. That's a fight about the structure of the whole, which is more fundamental (and more destabilizing) than fights within the structure.

The Paradox of Sovereignty

Sovereignty is the institution that claims to be the final institutional authority. The sovereign can make law but isn't bound by law (or is bound only by self-imposed constitutional limits). This creates a paradox: if sovereignty is the highest authority, what makes it legitimate? It can't be authorized by anything higher (that would be the sovereign). It must be self-authorizing.

The traditional answer is that sovereignty comes from "the people." The people are sovereign, and they delegate authority to government. This solves the legitimacy problem (popular sovereignty is self-evidently legitimate) but creates a new problem: the people are a multitude, not a unified agent. How do they exercise sovereignty?

Through representatives, institutions, constitutions, and elections. But all of these are creations of the sovereign

people. So the sovereign people create institutions to exercise their sovereignty, and the institutions define who counts as "the people" and how they can act. The sovereign is both creator and creation of the institutional structure.

This circularity is unavoidable. Institutions can't authorize themselves without already existing. They can't exist without being authorized. The circle is resolved through collective belief: enough people believe the institutions are legitimate, so they act as if the institutions have authority, and the institutions function as if they have authority, and the mutual belief creates the reality it assumes.

This makes sovereignty a particularly pure case of institutional reality: it exists because we collectively believe it exists and act accordingly. There's no metaphysical foundation. There's just the

circular structure of institutions creating themselves through collective recognition. The circle is vicious logically but virtuous practically. It works as long as enough people believe it works.

When belief fails (revolution, state collapse, delegitimization), sovereignty evaporates. The institutions claiming final authority lose that authority. The question “who is sovereign?” becomes open again. New claimants emerge. Eventually one establishes itself through some combination of force, legitimacy, and collective acceptance. Sovereignty reconstitutes itself, and the circular structure re-emerges.

The Upshot

States and institutions are abstract beasts: real entities with causal powers that exist primarily as patterns of behavior sustained by collective belief. They’re

composed of people but not reducible to people. They're structured by laws but not merely law. They're sustained by belief but not mere illusion.

The abstractness makes them strange but not less real. They employ people, extract resources, distribute goods, enforce rules, and make war. Their causal impact on the world is enormous. They shape your life more than most physical objects you encounter. The fact that they're abstract patterns rather than physical things doesn't make them less powerful.

What it does mean is that they're vulnerable in distinctive ways. Physical objects persist through structural integrity. Institutions persist through collective belief. Destroy the belief and the institution evaporates even if all the people and buildings remain. This makes institutions simultaneously more stable

(they can persist through complete replacement of parts) and more fragile (they can collapse if collective belief fails) than physical objects.

Understanding institutions as abstract wholes sustained by distributed belief helps explain their dynamics: why they're resistant to reform (each part resists changes that threaten it), why they sometimes collapse rapidly (collective belief fails), why their power often exceeds what their nominal structure would suggest (informal networks overlay formal hierarchy), and why responsibility for institutional actions is so hard to assign (distributed agency and emergent outcomes).

The practical lesson is that changing institutions requires changing beliefs, incentives, and structures simultaneously. Changing beliefs without changing structures creates cynicism (we know it

should work differently but it doesn't). Changing structures without changing beliefs creates resistance (formal rules change but informal practices persist). Changing incentives without changing beliefs or structures creates corruption (people game the new incentives in ways that undermine reform).

Successful institutional change requires operating at all three levels: cultivate new beliefs about legitimacy, redesign structures to embody those beliefs, and align incentives with the desired behavior. This is rare, which is why most reform fails and why institutions are so persistent even when they're dysfunctional.

States and institutions are the most powerful wholes humans create. They exist primarily as ideas instantiated in practice and sustained by collective belief. That makes them abstract, but not

weak. Ideas with billions of believers behind them are among the most powerful forces on Earth.

12. Minimalism vs Maximalism

How many objects are there in this room right now?
Count the furniture, the obvious answer is maybe ten or twenty. Count every distinct object (books, pens, dust particles), you're in the thousands or millions. Count every possible fusion of objects (book-plus-chair, left-half-of-desk-plus-ceiling-fan, all-objects-within-three-feet-of-the-window), you're at infinity.

Which count is correct? The answer reveals your ontological commitments about what exists. Minimalists say: only the fundamental stuff exists, everything else is conceptual bookkeeping. Maximalists say: every logically possible combination of fundamental stuff exists as a composite object. The middle contains dozens of positions, each with

arguments and problems.

This chapter is about the war over how much stuff there really is, and why the answer matters for everything else we've discussed. Because if arbitrary fusions exist, then the Ship of Theseus has infinite companions (ship-minus-one-plank, ship-plus-random-atoms, ship's-left-half), and most of our intuitions about objects are hopelessly confused.

How Many Things Exist, Really?

Start with the obvious answer: things exist if they're unified, bounded, and causally cohesive. Your desk exists. The random collection of your desk's left half plus some dust on Mars doesn't exist. This is common sense.

But common sense runs into problems immediately. A cloud is not unified (droplets constantly join and leave), not bounded (edges are fuzzy), and barely

causally cohesive (droplets interact weakly). Yet clouds exist. You can see them. They appear on weather radar. They cause rain. If clouds exist despite violating all three criteria, what are the criteria actually doing?

Maybe the answer is pragmatic: things exist if it's useful to track them. We track desks, clouds, and nations because tracking them helps us navigate the world. We don't track desk-plus-Mars-dust because that fusion serves no purpose. Existence is stance-dependent: relative to human purposes, desks exist and arbitrary fusions don't.

This is intellectually honest but metaphysically unsatisfying. It makes existence depend on human interests. Before humans existed, did clouds exist? Common sense says yes. Pragmatism says clouds existed as physical phenomena (water droplets suspended in

air) but not as objects (bounded, unified things) because there were no minds to treat them as objects. The physical stuff was there, but the object-carving is our contribution.

Or maybe the answer is scientific: things exist if they appear in our best scientific theories. Fundamental physics talks about quantum fields and particles. Higher-level sciences talk about molecules, cells, organisms, ecosystems. These are the real objects. Everything else is either reducible to these (supervenes on fundamental physics) or is useful fiction (nations, corporations, social roles).

This works better but still has problems. Our best scientific theories are incomplete and provisional. They'll change. Does that mean which objects exist changes when our theories change? That seems backwards. The objects are

supposed to ground our theories, not the other way around.

Each answer to "how many things exist?" presupposes a view about what makes something real versus merely conceptual. And the views are incompatible. You can't simultaneously believe that only fundamental stuff is real (minimalism), that all fusions of fundamental stuff are real (maximalism), and that reality is carved by human purposes (pragmatism). You have to pick.

Ontological Dieting vs Ontological Hoarding

Minimalism is ontological dieting: accept only what you absolutely need to explain the phenomena. This is intellectually austere and deeply aligned with scientific practice. Scientists constantly ask "can we explain this with fewer entities?" Every eliminated entity simplifies the

theory.

The extreme version is mereological nihilism: only fundamental simples exist, everything else is just simples arranged in patterns. There are no tables, only simples arranged table-wise. No organisms, only simples arranged organism-wise. The simples are real. The arrangements are real. The composite objects are conceptual overlays on arrangements.

This view has elegant simplicity. One kind of thing exists (simples), and everything supervenes on how they're arranged. No vagueness about composition (it doesn't occur), no worry about arbitrary fusions (they don't exist), no puzzle about persistence (only simples persist, arrangements change).

But the cost is that almost everything you care about doesn't exist. You don't exist (you're just simples arranged person-

wise). Your relationships don't exist. Your projects don't exist. Nations, corporations, trees, and clouds definitely don't exist. Only simples and their arrangements exist, and arrangements are not objects, just ways simples are configured.

This is pragmatically unlivable. You cannot think or act without treating composite objects as real. You can philosophically believe you're just simples arranged person-wise while practically treating yourself as a unified agent making choices. But the disjunction between philosophical belief and practical necessity is uncomfortable.

Maximalism is ontological hoarding: accept everything that could possibly exist. Every fusion of objects is itself an object. Every temporal stage of an object is an object. Every possible boundary-drawing creates an object. This is

intellectually promiscuous but eliminates vagueness. No borderline cases (everything exists), no puzzle about when composition occurs (always), no question about which fusions are real (all of them).

Unrestricted composition, discussed earlier, is the core maximalist commitment. It generates infinite objects in any region of space. Sharing the room with you right now are: you, your left half, your right half, you-minus-one-atom, you-plus-the-chair, you-during-this-second, you-during-the-next-second, and infinite others. They're all equally real. Most are useless to think about, but that's a pragmatic fact, not a metaphysical one.

This view has the virtue of clean principles. Composition always occurs. Every fusion exists. No special conditions required. The metaphysics is

simple even though the ontology is bloated.

But the cost is accepting infinite objects you'll never use or think about. Every region of space is infinitely crowded with overlapping objects. The desk is there, and so is desk-minus-one-molecule, and desk-plus-one-random-photon, and infinite others. They're all real in exactly the same sense. The one you call "the desk" is no more fundamental than any of the others. It's just more salient to you.

Why Both Sides Are Cheating

Minimalism cheats by treating simples as unproblematic. But what are simples? Atoms were thought to be simple (the name means "indivisible"), then we found they're composed of protons, neutrons, and electrons. Protons and neutrons are composed of quarks. Quarks might be composed of strings. Maybe it's

turtles all the way down.

If there are no fundamental simples, mereological nihilism collapses. You can't say "only simples exist" if there are no simples, only composites all the way down. The view requires a foundation it might not have.

Even if simples exist, it's not clear they're objects in any interesting sense. Quantum fields are the most fundamental entities in modern physics, and they're not particle-like objects. They're field configurations that can be excited in particle-like ways. Are field excitations "simples"? They don't have clear boundaries or determinate locations. They're not little billiard balls sitting in space. Treating them as paradigm objects and everything else as derivative seems backwards.

Minimalism also cheats by treating "arranged F-wise" as unproblematic.

What does it mean for simples to be arranged table-wise? Just that they're in the configuration that would compose a table if composition occurred. But this smuggles composition back in. You're saying there's no table, just simples configured table-ishly, where "table-ish configuration" means "the configuration that would make a table." This is either circular or it presupposes you know what tables are independent of composition, which undermines the whole point.

Maximalism cheats by treating existence as cheap. If everything exists, "existence" stops doing work. Normally, saying "X exists" distinguishes X from things that don't exist (unicorns, round squares, Santa Claus). But if unrestricted composition is true, almost every describable fusion exists. The fusion of all unicorn-shaped arrangements of atoms exists (though it might be scattered

and useless). The fusion you-minus-your-left-pinky exists right now.

This makes “exists” almost meaningless as a discriminator. The real work is done by salience: which existing things matter for explanation or action. But then we’re back to pragmatism. The metaphysics says everything exists. The practice says most existing things don’t matter. The metaphysics and practice are talking past each other.

Maximalism also cheats about identity. If infinite overlapping objects occupy every region, identity claims become bizarre. “This desk” could refer to any of infinite desk-like fusions. Which one are you pointing at? The answer can’t be “the one I mean” because your meaning doesn’t select among metaphysically equivalent fusions. They all have equal claim to being “the desk.” Maximalism creates massive indeterminacy in reference that

it doesn’t have the resources to resolve.

Choosing Your Metaphysics Like a Toolset

Maybe the right approach is instrumentalism: don’t ask what really exists, ask which ontological framework is useful for which purposes.

For fundamental physics, minimalism works well. Physics wants to reduce explanation to the simplest possible entities and laws. Whether composites “really exist” is irrelevant. What matters is whether you can explain their behavior in terms of fundamental constituents. If you can, the composites add nothing to the explanation.

For biology, moderate compositional realism works better. Organisms are real. They’re not just particles arranged organism-wise. They’re functionally integrated systems with emergent

properties that matter for biological explanation. Natural selection operates on organisms, not on particles. Treating organisms as real captures something important about the explanatory structure.

For sociology, institutional realism is essential. You cannot explain social phenomena without treating institutions as real entities with causal powers. The corporation does things. The nation does things. Reducing these to individuals misses system-level dynamics that drive the explanation.

On this view, ontology is plural. Different domains require different object-carving because different explanatory structures privilege different levels of organization. The question isn't "what really exists?" but "what needs to exist for this explanation to work?"

This is pragmatism with a sophisticated

face. It avoids dogmatism (there's no one true ontology) and pluralism (different ontologies for different purposes). But it still makes existence relative to explanatory frameworks, which is metaphysically uncomfortable.

Are organisms real independent of biological explanation, or are they only real relative to biological explanation? The first seems right (organisms were real before biology existed), but it commits you to some ontology being the true one. The second respects explanatory pluralism but makes existence framework-relative in ways that seem wrong.

The Virtues of Ontological Modesty

Here's a different approach: accept that we don't and maybe can't know what really exists. Our access to reality is always mediated by concepts, language, and purposes. We carve the world in

ways that serve our needs. The carvings correspond to real patterns in nature (we're not inventing the patterns), but which patterns constitute objects is partly our contribution.

This is ontological modesty: acknowledge the limits of our ability to answer metaphysical questions definitively. We know there's stuff. We know it has structure. We know some carvings are better than others (they track natural boundaries, enable prediction, support action). But we don't know whether "object" is a fundamental metaphysical category or a pragmatic one we impose.

The virtue of modesty is that it lets you use different ontologies for different purposes without claiming one is the true ontology. You can think in terms of people when discussing moral responsibility, systems when analyzing

institutional failures, particles when doing physics. The switch between ontologies is not a metaphysical error but a pragmatic necessity.

The vice of modesty is that it gives up on answering the questions. It says “we can’t know” when you want to know. It leaves you with practical ontologies that work for various purposes but no theoretical ontology that’s ultimately true. This is intellectually modest but existentially unsatisfying.

Where the Debate Leads

The minimalism-maximalism debate has been running in philosophy for decades without resolution. This suggests the debate might be intractable: the considerations on each side are strong enough to support their view but not strong enough to defeat the other side.

Minimalism is right that ontological

parsimony is a theoretical virtue. Don't multiply entities beyond necessity. If you can explain everything with just simples arranged in patterns, you should. Simpler theories are better, all else equal.

Maximalism is right that vagueness is a theoretical vice. Don't accept theories with arbitrary cutoffs or borderline cases. If you admit that some things compose, you need principled criteria for when composition occurs, and no proposed criteria avoid vagueness or arbitrariness. Unrestricted composition is the only vagueness-free answer.

Pragmatism is right that our ontological commitments are shaped by our explanatory and practical needs. We carve reality in ways that serve our purposes. Different purposes privilege different carvings. This doesn't make the carvings arbitrary (they have to track real patterns), but it makes them purpose-

relative.

All three views capture something important. None defeats the others. The debate continues because the values in tension (parsimony, precision, pragmatic utility) can't be fully satisfied simultaneously. You can have parsimony with vagueness (minimalism), precision with bloat (maximalism), or utility with relativity (pragmatism). You can't have all three.

The Practical Upshot

For practical purposes, the debate doesn't matter much. Whether or not arbitrary fusions "really exist," you're not going to treat desk-plus-Mars-dust as a salient object. Whether or not composites "really exist" or are just simples arranged, you're going to treat your body as a unified object, not as a pattern of simples.

The practical ontology is driven by salience, not metaphysics. You track objects that matter for prediction, explanation, and action. The metaphysical question of whether those objects are “really real” or “just useful fictions” doesn’t affect how you treat them.

Where the debate does matter is in understanding what you’re committed to when you make mereological claims. If you say “the nation did X,” are you claiming nations are real entities with causal powers, or are you using shorthand for a complicated claim about patterns of individual behavior? The answer affects how you think about responsibility, change, and intervention.

Similarly, if you say “this is the same ship,” are you claiming identity is metaphysically determinate, or are you claiming that for your purposes (legal

ownership, historical connection, functional continuity), treating it as the same ship makes sense? The answer affects what arguments can resolve identity disputes.

The metaphysical debate clarifies what work your claims are doing and what they're committed to. It doesn't settle the practical questions, but it reveals what's at stake in how you answer them.

The Takeaway

Minimalism says almost nothing exists (just simples). Maximalism says almost everything exists (all fusions). Pragmatism says existence is relative to purposes. They're incompatible, none is clearly right, and the debate is probably intractable.

The reason to care is not that you need to pick a side. It's that understanding the options clarifies what you're assuming

when you talk about objects. Are you committed to objects being fundamental or derivative? Are you allowing vagueness or requiring precision? Are you treating existence as absolute or purpose-relative?

These commitments shape everything else: what counts as persistence through change, when composition succeeds or fails, whether wholes are reducible to parts, whether arbitrary fusions are real. The whole edifice of mereological thinking rests on these foundations, and the foundations are contested.

The philosophically mature position is probably ontological pluralism tempered with modesty. Use the ontology that fits your explanatory purposes. Recognize that different purposes privilege different ontologies. Don't claim your preferred ontology is the uniquely correct one. Accept that we're carving a world rich

enough to support multiple valid carvings, and the choice between carvings is partly ours.

This won't satisfy metaphysical realists who want to know what really exists. But it's honest about the limits of our ability to answer that question definitively. And it's pragmatically workable: you can navigate the world, build theories, and make moral judgments without settling whether simples or composites or both or neither are "really real."

The world doesn't care about our ontological commitments. It just keeps being complex, supporting multiple decompositions, resisting our attempts to impose one true carving. The mature response is to accept the plurality and work with it rather than pretending there's a fact of the matter we're discovering.

13. Are Wholes Ever Fundamental?

Water is two hydrogen atoms and one oxygen atom. Full stop. If you know everything about hydrogen and oxygen, you know everything about water. Water is nothing over and above its atomic constituents. Reductionism wins.

Except water is wet, and hydrogen and oxygen aren't wet. Water has a freezing point, neither hydrogen nor oxygen have that freezing point. Water dissolves salt, the individual atoms don't. These properties are not properties of the atoms. They're properties of the configuration, the pattern, the whole.

So maybe water is not just hydrogen and oxygen. Maybe the whole has properties the parts lack, properties that can't be predicted from studying parts in isolation. Maybe wholes sometimes push back, constraining their parts in ways that

make the whole more fundamental than reductionism suggests.

This chapter is about the uncomfortable possibility that parts might not be all there is, that wholes might do genuine work in the universe, and that emergence might be real without being mystical.

Minimalism vs Maximalism (Revisited From Below)

The previous chapter approached ontology from above: how many objects exist? This chapter approaches it from below: which objects are fundamental?

Reductionism says parts are fundamental. Wholes supervene on parts (you can't change the whole without changing the parts). Wholes are determined by parts (specify all facts about parts and their arrangement, and you've specified everything about the whole). Therefore wholes are ontologically dependent on

parts. The parts are doing the real work.

This seems obviously true. Your body is made of cells. Cells are made of molecules. Molecules are made of atoms. Atoms are made of particles. At each level, the higher depends on the lower. The dependency flows one direction: down. Change the particles and you change everything above. Change "the organism" without changing any particles and… what does that even mean? You can't. Wholes depend on parts, not the other way around.

But dependency isn't the same as fundamentality. The whole might depend on the parts for its existence while still being explanatorily or causally fundamental. The organism is a whole that emerges from cells, but once it exists, it constrains what its cells can do. The constraint is top-down causation: the whole affecting the parts.

Or consider: the economy depends on individual transactions. You can't change the economy without changing some transactions. But the economy also creates incentives that shape individual behavior. High unemployment discourages job-seeking (why try when jobs are scarce?). Interest rates affect borrowing decisions. Market prices coordinate behavior. These are whole-level properties affecting part-level actions.

This suggests a different picture: reality has multiple levels, each with its own dynamics. Lower levels constrain higher levels (you need molecules to have organisms). Higher levels constrain lower levels (organismal needs shape cellular behavior). The constraint is bidirectional. Neither level is simply fundamental. Both are doing causal work.

Emergence Without Woo

“Emergence” is a suspect word in philosophy. It gets used to handwave away hard problems. “Consciousness emerges from brain activity” sounds like an explanation but might just be restating the problem. How does it emerge? What is the mechanism? If you can’t answer, you’re not explaining, you’re labeling ignorance.

But there are clear cases of emergence that aren’t mysterious. Temperature is an emergent property of molecular motion. No individual molecule has a temperature. Temperature is a statistical property of large ensembles: the average kinetic energy of many molecules. You cannot predict temperature from studying one molecule. You need to know about many molecules and their distribution.

This is weak emergence: higher-level properties that are unexpected or

unpredictable from lower-level properties but that are still fully determined by them. Temperature emerges from molecular motion, but it's still just molecular motion viewed at a different scale. There's nothing extra, no downward causation, no autonomous dynamics at the temperature level.

Weak emergence is everywhere and uncontroversial. Liquidity emerges from molecular interactions (individual molecules aren't liquid). Wetness emerges from bulk properties. Conductivity, tensile strength, color—all emergent properties of materials that individual atoms lack.

The interesting question is whether there's strong emergence: higher-level properties with autonomous causal powers, properties that aren't just lower-level properties viewed differently but that genuinely constrain lower-level

behavior in ways that can't be fully explained by lower-level laws.

If strong emergence exists, wholes are doing real work that parts alone can't explain. The whole would be fundamental in a sense that challenges reductionism. The universe would have genuine levels, each with its own causal structure, rather than one fundamental level (particles) with derivative levels above it.

Strong vs Weak Emergence

The test for strong emergence is whether you need to invoke the higher level to explain lower-level behavior. If yes, the emergence is strong. If no, it's weak.

Temperature is weak emergence because you never need to invoke temperature to explain molecular motion. Molecules move according to physical laws (collision dynamics, forces, energy

conservation). Temperature is a useful summary of their aggregate behavior, but it doesn't cause anything. It's caused by molecular motion, it doesn't cause molecular motion.

But consider biology. Do organisms constrain cellular behavior in ways that can't be fully explained by cellular-level laws? Maybe yes. Cells in your liver behave differently than the same cells would behave in isolation because they're embedded in an organism that regulates them. The regulation is top-down: organismal needs (maintain blood sugar, filter toxins) shape cellular activity through hormones, neural signals, and mechanical constraints.

You might say this is still just lower-level causation. The hormones are molecules. The neural signals are electrochemical. The mechanical constraints are physical forces. Nothing

top-down is happening, just complex bottom-up interactions.

But the pattern matters. The hormones are produced by other organs in response to organismal needs. The neural signals coordinate behavior across distant body parts. The mechanical constraints are imposed by organismal structure. The organism as a whole creates an environment for its cells that wouldn't exist if the cells were independent. That environment shapes cellular behavior.

Is this strong emergence? It depends on whether "organismal needs" are reducible to cellular-level facts. If organismal needs are just patterns in cellular behavior, then explaining cellular behavior by organismal needs is circular or redundant. But if organismal needs are genuinely higher-level properties that can't be defined in cellular terms, then citing them in explanations adds

something the cellular level can't provide.

The debate here is fierce and unresolved. Reductionists say everything about organisms ultimately reduces to molecular biology, which reduces to chemistry, which reduces to physics. Emergentists say this reduction is impossible in practice and maybe in principle, and that higher-level laws are autonomous.

When Wholes Start Pushing Back

The clearest cases of wholes constraining parts are self-organizing systems. These are systems where local interactions between parts create global patterns, and the global patterns then constrain local interactions. The constraint is a feedback loop: parts create wholes, wholes shape parts.

Flocks are an example. Birds following

simple local rules (match velocity, maintain spacing, avoid predators) create flock-level patterns. The patterns (flock shape, direction, density) then influence individual birds (a bird on the edge might move toward the center, changing its local behavior in response to flock-level properties).

Is the flock pushing back on the birds? In one sense no—the flock is just the birds. In another sense yes—the flock has properties (shape, direction) that individual birds respond to, and those properties aren't properties of any individual bird. They're properties of the configuration. Birds are responding to a whole-level property that emerges from but isn't reducible to individual birds.

Markets are similar. Individual trades create market prices. Market prices then constrain individual trading decisions (you won't buy at a price above market,

won’t sell below it). The price is a whole-level property that emerges from individual actions but then constrains those very actions. The feedback creates market dynamics that can’t be explained purely from individual psychology. You need market-level concepts (liquidity, volatility, price discovery) to explain what happens.

Societies are the strongest case. Cultural norms emerge from individual behavior. But once established, norms constrain individual behavior (you’re punished for violations, rewarded for compliance). Language emerges from communication practices. But once established, language shapes thought (you can think certain things more easily in some languages than others). Institutions emerge from coordination. But once established, they channel behavior in ways individuals can’t easily resist.

These are all cases where wholes exert causal influence on parts. The influence is not mysterious—it operates through mechanisms (social pressure, feedback loops, information flows). But it's genuinely top-down: the whole affects the parts in ways that require whole-level concepts to explain.

Holism as Conceptual Necessity

Even if you don't believe in strong emergence ontologically, you might need it epistemologically. Some explanations require whole-level concepts that can't be reduced to part-level concepts without losing explanatory power.

Economics is the paradigm case. You can in principle reduce all economic phenomena to individual psychology and decision-making. Every transaction is individuals making choices. But actually doing the reduction is impossible. You cannot explain inflation, unemployment,

or market crashes by analyzing individual psychology. The explanations require aggregate concepts (money supply, labor markets, systemic risk) that operate at the whole level.

Trying to reduce economics to psychology is like trying to predict weather by modeling individual air molecules. It's possible in principle (weather supervenes on molecules) and impossible in practice (the computation is intractable). You need higher-level concepts (pressure systems, temperature gradients, humidity) to make predictions. The concepts refer to aggregate properties, not individual molecules.

This is methodological holism: the thesis that some explanations essentially require whole-level concepts. You cannot eliminate the whole-level concepts without losing the ability to explain. This doesn't mean wholes are ontologically

fundamental (maybe they're still just parts arranged). But it means they're explanatorily indispensable.

Biology is increasingly methodological holist. Molecular biology tried to reduce everything to genes and proteins. But organisms turned out to be too complex for pure bottom-up explanation. You need concepts like "developmental program," "regulatory network," and "organism-environment interaction" that operate at the organismal or systems level. These concepts might ultimately supervene on molecular facts, but you can't do biology by just studying molecules.

Neuroscience faces the same issue. Consciousness might supervene on neural activity, but explaining consciousness requires concepts about information integration, global workspace, attention, and self-modeling

that don't obviously reduce to neural firing patterns. Maybe they do reduce (eventually), but neuroscience can't wait for that reduction. It needs whole-level concepts now to make progress.

The epistemological point is that even if reductionism is true ontologically (wholes are nothing but parts), it fails as a research program. Science requires multiple levels of explanation. Each level has its own concepts, laws, and methods. Trying to do all science at the fundamental physics level would paralyze research. You need autonomous higher-level sciences.

The Uncomfortable Possibility That Parts Are Secondary

Now the truly disturbing thought: what if parts are the derivative level and wholes are fundamental?

Quantum mechanics suggests this might

be true. Particles don't have determinate properties until measured. What exists before measurement is a wave function describing the system as a whole. The wave function is holistic: it describes correlations between parts (entanglement) that can't be captured by describing parts independently.

When you measure a particle, the wave function "collapses" and the particle has a definite property. But the collapse is a whole-system event, not a part-level event. The whole system (particle plus measurement apparatus plus environment) collapses into a definite state. The particle's property is determined by the whole-system collapse, not by anything intrinsic to the particle.

This suggests particles are not fundamental. The quantum state is fundamental, and particles are

abstractions from the quantum state that emerge when we perform measurements. The whole (quantum state) is primary. The parts (particles) are secondary, derived from interactions with measurement apparatus.

Most physicists don't talk this way because it sounds weird and mystical. But the mathematics of quantum mechanics points in this direction. Entangled states are the norm. Product states (where parts are independent) are special cases. The holistic level is fundamental, the particulate level is derived.

This is deeply counterintuitive. Our everyday experience is of discrete objects (parts) that combine into wholes. But quantum mechanics suggests that everyday experience is emergent. The fundamental level is fields, not particles. The fields are holistic, not decomposable

into independent parts. Parts emerge through our interaction with the fields.

Organizational Closure and Self-Maintenance

Another angle on whole-level fundamentality comes from organizational closure: systems that maintain themselves through circular causal processes where each part depends on other parts for its existence.

Organisms are organizationally closed. Your cells produce proteins that regulate cellular function. Cellular function maintains tissues. Tissues compose organs. Organs coordinate to keep the organism alive. The organism being alive is necessary for cells to function. The loop is: cells → tissues → organs → organism → cells. Each part depends on the whole, and the whole depends on the parts.

This creates a strange ontological status. The organism is not just a collection of cells because the cells depend on the organism for their maintenance. But the organism is not separate from its cells because it's composed of them. The organism and its cells are mutually dependent in a way that makes neither level clearly more fundamental.

Autopoiesis is the term for this self-maintaining organization. Autopoietic systems produce the components that produce them. They're circular: the process generates the structure, and the structure enables the process. This circularity makes it unclear whether parts or wholes are primary.

Cells are autopoietic. They produce proteins and lipids that form membranes. The membranes enclose the processes that produce proteins and lipids. The cell maintains itself through this loop. If you

break the loop (stop protein production), the cell dies and decomposes into parts that no longer constitute a cell.

Organisms are autopoietic. They metabolize to maintain their structure. The structure enables metabolism. Break the loop (death) and you have a corpse: parts that used to compose an organism but no longer do because the organizational closure is broken.

Societies might be autopoietic. They produce institutions that organize behavior. Organized behavior sustains institutions. The loop maintains society. Revolutionary collapse is breaking organizational closure: the institutions stop organizing behavior, behavior stops sustaining institutions, and society fragments into groups that no longer compose a unified whole.

Autopoietic systems suggest that wholes are not just collections of parts. They're

processes of self-maintenance where the whole and parts mutually constitute each other. You can't say the parts are fundamental because the parts depend on the whole for their identity and function. You can't say the whole is fundamental because it's composed of parts. Both levels are equally real and mutually dependent.

The Upshot

Are wholes ever fundamental? The answer is: sometimes yes, depending on what you mean by "fundamental."

If "fundamental" means "ontologically prior," then parts are probably fundamental (at least in classical physics). Wholes depend on parts for their existence. Change the parts and you change the whole. The dependency flows bottom-up.

If "fundamental" means "explanatorily

indispensable," then wholes are often fundamental. Many explanations require whole-level concepts that can't be eliminated without losing explanatory power. Science needs multiple levels. Each level is fundamental for explanations at that level.

If "fundamental" means "causally autonomous," then some wholes are fundamental. Organizationally closed systems, self-organizing systems, and entangled quantum systems all exhibit top-down causation where wholes constrain parts in ways that require whole-level concepts to explain.

The reductionist picture (parts all the way down, wholes are just bookkeeping) is too simple. Reality has genuine levels. Higher levels emerge from lower levels but aren't fully reducible to them. The levels interact bidirectionally. Parts constrain wholes (you can't have

organisms without cells). Wholes constrain parts (organisms regulate cells). Neither level is simply fundamental.

This is ontological pluralism again, but now applied to levels rather than objects. Just as different purposes privilege different object-carvings, different explanatory projects privilege different levels. Physics works at the particle level. Biology works at the organismal level. Economics works at the market level. Each level is fundamental for its domain.

The mature view is that fundamentality is relative to explanatory context. Parts are fundamental for some purposes. Wholes are fundamental for others. Asking which is really fundamental, independent of any context, is asking a question that might not have an answer.

This won't satisfy metaphysical realists who want to know the ultimate furniture

of reality. But it's honest about the complexity. The world has structure at multiple scales. Each scale has its own dynamics. The scales interact. No single scale is uniquely fundamental. That's not a failure of our understanding. It's a feature of a world complex enough to support multiple levels of organization, each doing real work.

Wholes push back. Sometimes gently (weak emergence), sometimes forcefully (strong emergence, organizational closure, quantum holism). The pushing is not mystical. It's mechanistic, explicable, and essential to understanding how complex systems work. Once you see it, reductionism starts to look like a useful approximation that works for simple systems and fails for complex ones. And most of what matters is complex.

14. Living in a World of Leaky Boundaries

You are reading this sentence. But what, exactly, is doing the reading?
Your eyes are scanning symbols. Your visual cortex is processing patterns. Your language centers are extracting meaning. Your memory is integrating this with past knowledge. Your attention is (maybe) focused on comprehension. Thousands of neural processes are firing simultaneously, none of which is "you" in isolation, all of which compose into the reading experience.

Where does "you" begin and end in this process? Are you the neural activity? The experience of reading? The body housing the brain? The person with a history and future who chose to read this? The answer shifts depending on what you're trying to explain. And that's the point:

boundaries are not discovered, they're drawn. And they leak constantly.

This final chapter is about what you've learned by following mereology into its strange corners, and what it means to navigate a world where every object you encounter is actually a negotiated fiction with porous boundaries and context-dependent identity.

Why Clean Categories Are a Fantasy

You want reality to come pre-divided into discrete units. You want "tree" to be a natural kind with clear boundaries: if something is a tree, it's fully a tree, and if it's not a tree, it's fully not a tree. No borderline cases, no vagueness, no context-dependence.

This want is understandable and completely unsatisfied by reality.

Is a bonsai a tree? It has the right biological characteristics (woody stem,

perennial growth, branches). But it's three inches tall. Trees are supposed to be… tall? How tall? Where's the cutoff? There isn't one. "Tree" is a vague category with prototypical members (oak, pine, redwood) and borderline cases (bonsai, large shrubs, tree-like plants).

Is a virus alive? It reproduces (with help from cells). It evolves. It has genetic material. But it's not cellular. It can't metabolize independently. It exists as crystalline structures when not infecting hosts. The question "is it alive?" doesn't have a determinate answer because "alive" is a cluster concept: typical life has all these features, but some entities have only some of them. Viruses are borderline alive.

Is a corporation a person? Legally, yes (it can own property, enter contracts, sue and be sued). Metaphysically, no (it's an

organizational structure, not a biological organism). Morally, complicated (it has interests and makes decisions, but it lacks consciousness and phenomenal experience). The question mixes different senses of "person" that don't align.

Clean categories are fantasies we impose on messy reality. Reality has clusters, prototypes, and family resemblances. It has clear cases and borderline cases. It has objects that are trees for some purposes (biological classification) and not trees for other purposes (legal definitions of forest land). The categories are ours, serving our interests, and reality tolerates many different categorization schemes simultaneously.

This is not relativism. Some categorizations are better than others (they carve at natural joints, track causal structure, enable prediction). But "better" is still relative to purposes. Biological

classification is better for biology. Legal definition is better for law. There's no purpose-independent "best" categorization that captures what things really are.

How Mereology Trains You to Spot Bullshit

Once you've worked through mereological problems, you develop a nose for conceptual confusion disguised as factual disagreement.

Someone says: "The problem is individual responsibility has declined. People don't take ownership of their actions anymore. They blame the system."

Someone else says: "The problem is systemic. Structures create outcomes individuals can't control. Focusing on individual responsibility is victim-blaming."

This looks like a factual disagreement. But it's often a mereological disagreement: should we analyze at the individual level or the systemic level? Both levels are real. Both do causal work. The disagreement is about which level is more important for this context, not about the facts.

Mereology teaches you to ask: "What level of analysis are we using, and is that level appropriate for what we're trying to explain?" Sometimes the answer is individual (personal choice matters). Sometimes systemic (structural constraints dominate). Often both (individuals make choices within structures, structures are sustained by individual behavior). The confusion comes from thinking you have to pick one.

Or someone says: "Corporations are just people. They don't have interests

separate from their members. Blaming corporations is just avoiding individual accountability."

Someone else says: "Corporations are distinct entities with emergent properties. They act in ways individuals don't intend. They need to be regulated as wholes."

Mereological analysis reveals this is about whether to treat corporations as reducible to members (methodological individualism) or as autonomous wholes (methodological holism). Both frames are valid for different purposes. Legal liability works better with holism (sue the corporation). Moral responsibility works better with individualism (individuals make decisions). The disagreement is about which frame to use, not about metaphysical facts.

Or someone says: "There's no objective way to define a nation. National

boundaries are social constructs that change over time. Nations are imagined communities."

Someone else says: "Nations are real. They have governments, territories, and populations. They do things. Social construction doesn't make them less real."

Mereology reveals both are right. Nations are socially constructed (their boundaries are conventional, sustained by collective recognition). Nations are real (they have causal powers, they constrain behavior, they persist through changes). Social construction and reality are not opposites. Many real things are socially constructed. The question is what kind of real they are (institutional versus physical), not whether they're real.

The pattern is: apparent factual disagreements often hide mereological

disagreements about levels, boundaries, or modes of composition. Spotting the underlying mereology lets you dissolve fake disagreements and focus on real ones.

Thinking in Seams Instead of Objects

The traditional way to think about reality is object-oriented. There are things. Things have properties. Things interact. Understand the things and you understand everything.

Mereology suggests a different orientation: think in terms of seams. Where do objects leak into each other? Where do boundaries blur? Where does the environment start to do work that looks like part-level work? Where do wholes constrain parts and parts rebel against wholes?

Seams are where interesting things happen. They're where composition

succeeds or fails. They're where responsibility leaks. They're where different systems interface and create emergent behavior.

Your skin is a seam between you and environment. But it's not a clean boundary. Bacteria live on it. Chemicals pass through it. Temperature affects it. Damage it and the boundary fails (infection). The skin is less a wall and more a negotiated interface that mostly succeeds at maintaining distinction but constantly threatens to fail.

The cell membrane is a seam between cell and environment. It's selectively permeable: some things pass freely, others are blocked, others are actively transported. The selectivity is what makes the cell a cell. Destroy the membrane and the cell dissolves into its environment. The boundary does work. It's not just where the cell ends but what

maintains the cell's existence.

The corporation-employee relationship is a seam. Employees are not parts of the corporation (they can leave). But they're not fully independent (they're bound by contracts and loyalty obligations). The relationship is contractual, voluntary, and asymmetric (corporations have more power). The seam is where individual agency meets institutional structure. Understanding corporations requires understanding this interface, not just the corporation or employees in isolation.

Ecological boundaries are seams. Where does the forest end and the grassland begin? There's usually a transition zone where trees thin out, grasses increase, and species from both ecosystems mix. The boundary is not a line but a gradient, a seam where different systems blend. The interesting ecology happens at these edges, where organisms from different

systems interact.

Thinking in seams means attending to interfaces, transitions, boundaries, and the places where systems leak into each other. This is often more productive than thinking about objects in isolation because the leakage is where the action is.

The Final Insult: You Are Both a Part and a Temporary Whole

Throughout this book, you've been treated as a reader, a person, a unified agent making choices about whether to keep reading. This is a useful fiction, and it's mostly true. But mereology insists on the final uncomfortable truth: you are not one thing.

You are a part of larger wholes. You're part of a family, a workplace, a community, a nation, an ecosystem. These wholes shape your behavior

through norms, incentives, relationships, and material constraints. You're not fully autonomous. Your choices are partly determined by the wholes you're embedded in.

You are a temporary whole composed of parts. Your cells turn over. Your molecules are replaced. Your beliefs change. Your memories fade. Your personality shifts. The parts composing you are constantly changing, and yet "you" persist through the change. But the persistence is approximate, not absolute. Childhood you and current you share some continuity (memory, bodily continuity, causal connection) but are also profoundly different.

You are both simultaneously. You're a whole relative to your parts (cells, organs, thoughts). You're a part relative to wholes you're embedded in (relationships, institutions, societies).

Neither description is complete. Both are necessary.

This dual nature creates tension. As a whole, you want to maintain your integrity, your boundaries, your autonomy. You want to persist through change without dissolving into parts. As a part, you're subject to forces from larger wholes that you can't fully control or escape. You want to integrate into communities, relationships, and institutions, but integration means surrendering some autonomy.

The tension is not resolvable. You are fundamentally both part and whole, and the demands of these roles conflict. The best you can do is navigate the tension: maintain enough integrity to be a coherent agent, maintain enough integration to participate in larger wholes. Too much integrity and you're isolated. Too much integration and you

dissolve into roles and relationships with no coherent self.

And it's temporary. The whole that is you exists for a human lifespan, maybe eighty years if you're lucky. Before birth, the parts existed but the whole didn't. After death, the parts will persist (as atoms, as nutrients recycled into ecosystems, as effects on others) but the whole will be gone. You are a pattern in matter that has temporary stability, not permanent existence.

This is not nihilism. Temporary wholes are still real. They do things. They matter. But they're contingent, not necessary. The universe doesn't care whether you exist. You exist because the right parts came together in the right configuration at the right time. Change the configuration and you cease to exist, even if all the parts remain.

What Mereology Gives You

If you've followed this book carefully, you've acquired a set of tools for thinking about composition, persistence, boundaries, and levels. The tools don't give you final answers (there often aren't any). They give you better questions.

Instead of "what is this thing?" ask "what boundaries am I drawing and why?"

Instead of "is this one object or many?" ask "at what level should I analyze this?"

Instead of "who is responsible?" ask "how is responsibility distributed across levels and parts?"

Instead of "is this real?" ask "what kind of real is it—physical, institutional, social, conceptual?"

Instead of "what's the cause?" ask "what level of causation matters for this explanation?"

These questions don't have algorithmic

answers. They require judgment, context-sensitivity, and tolerance for ambiguity. But they're the right questions for navigating a world where objects have fuzzy boundaries, wholes emerge from parts in complex ways, and responsibility leaks between levels.

Mereology also gives you a kind of intellectual humility. Once you've seen how many mereological puzzles are intractable (the Ship of Theseus, arbitrary objects, personal identity through change), you stop expecting clean answers to questions about what things really are. You accept that reality is messier than our categories, that objects are partly conventional, that existence comes in degrees and types.

This humility is productive. It prevents you from mistaking your preferred carving of reality for the one true carving. It makes you attentive to how

other people are carving reality differently for their purposes. It lets you switch between levels and frameworks fluidly rather than insisting on one correct view.

And it gives you a vocabulary for talking about composition that's more sophisticated than folk mereology. Instead of vague talk about "parts and wholes," you can distinguish contact composition from functional integration from coordination, you can track boundary failures and compositional seams, you can identify when problems are part-level versus whole-level, you can see where responsibility is distributed versus concentrated.

The vocabulary doesn't solve problems by itself. But it lets you articulate problems precisely, which is the first step toward solutions. And it lets you communicate with others about

mereological issues without talking past each other.

Final Thoughts on Leaky Boundaries

The world is not made of discrete objects with clean edges. It's made of matter and energy organized into patterns, where the patterns have varying degrees of stability, integration, and autonomy. We carve the patterns into objects because objects are how human cognition works. We need discrete categories, clear boundaries, stable identities. The world gives us enough structure to support these carvings, but it doesn't enforce them.

This means every object you interact with has leaky boundaries. Your body leaks into your environment (sweat, breath, shed cells). Your environment leaks into your body (air, food, chemicals). Your thoughts leak into your actions. Your actions leak into social

effects. Your identity leaks across time (you're not exactly the same person you were, but you're not completely different either).

The leakage is not a bug. It's a feature of a world where composition is real, where parts interact, where wholes emerge from parts and constrain them, where boundaries are maintained by ongoing processes rather than fixed by metaphysical facts.

Living in this world means accepting the leakage. You cannot have perfectly bounded objects because perfect boundaries would require perfect isolation, and perfect isolation would mean no composition, no interaction, no emergent properties, no life. The leakage is what makes complexity possible.

It also means accepting that many questions don't have determinate answers. Is this the same ship? Depends

on what you care about (material continuity, functional continuity, historical continuity). Is the corporation responsible? Depends on whether you're assigning legal liability or moral guilt. Is consciousness a part of the brain or a whole-level property? Depends on your explanatory framework and what you're trying to predict.

The indeterminacy is not failure. It's the world being rich enough to support multiple true descriptions that emphasize different features. The descriptions are all correct, just answering different questions or serving different purposes.

Mereology is the study of this richness. It reveals how parts compose into wholes, how wholes constrain parts, how boundaries are drawn and maintained, how composition succeeds and fails, how objects persist through change, and how responsibility distributes across levels. It

doesn't give you a simple picture of reality because reality isn't simple. It gives you tools for navigating complexity without getting lost.

The final lesson is this: reality has joints, but they're not where you think they are. Every clean boundary is actually a seam. Every discrete object is actually a pattern with fuzzy edges. Every level is interconnected with other levels. The world is one continuous process with multiple stable patterns at different scales, and we call those patterns "objects" when it's useful to do so.

This is not a recipe for confusion. It's an invitation to think more carefully about what you're assuming when you talk about things. Parts exist. Wholes exist. Boundaries exist. But they're all negotiated, maintained, and context-dependent in ways that folk metaphysics doesn't acknowledge.

Once you see this, you can't unsee it. You'll notice the mereological structure of every argument, every explanation, every object, and every institution you encounter. You'll see where boundaries are doing work and where they're arbitrarily imposed. You'll see where composition is robust and where it's fragile. You'll see where wholes are genuinely autonomous and where they're just convenient bookkeeping.

This is what mereology gives you: clear sight into the structure of reality, unobscured by the fantasy that objects are simple, discrete, and cleanly bounded. They're not. They're complex, leaky, and contingent. And that's exactly what makes them interesting.

Welcome to a world where everything is made of something else, where wholes emerge and dissolve, where boundaries shift, and where the question "what is

this really?” rarely has one correct answer. It’s messier than you wanted. It’s richer than you imagined. And now you have the tools to navigate it.

Marcus Qadar